



US 20150039859A1

(19) **United States**

(12) **Patent Application Publication**
Abdallah

(10) **Pub. No.: US 2015/0039859 A1**

(43) **Pub. Date: Feb. 5, 2015**

(54) **MICROPROCESSOR ACCELERATED CODE OPTIMIZER**

(52) **U.S. Cl.**

CPC **G06F 9/38** (2013.01); **G06F 9/30145** (2013.01)

(76) Inventor: **Mohammad Abdallah**, San Jose, CA (US)

USPC **712/206**

(21) Appl. No.: **14/360,282**

(22) PCT Filed: **Nov. 22, 2011**

(86) PCT No.: **PCT/US2011/061957**

§ 371 (c)(1),
(2), (4) Date: **Sep. 15, 2014**

Publication Classification

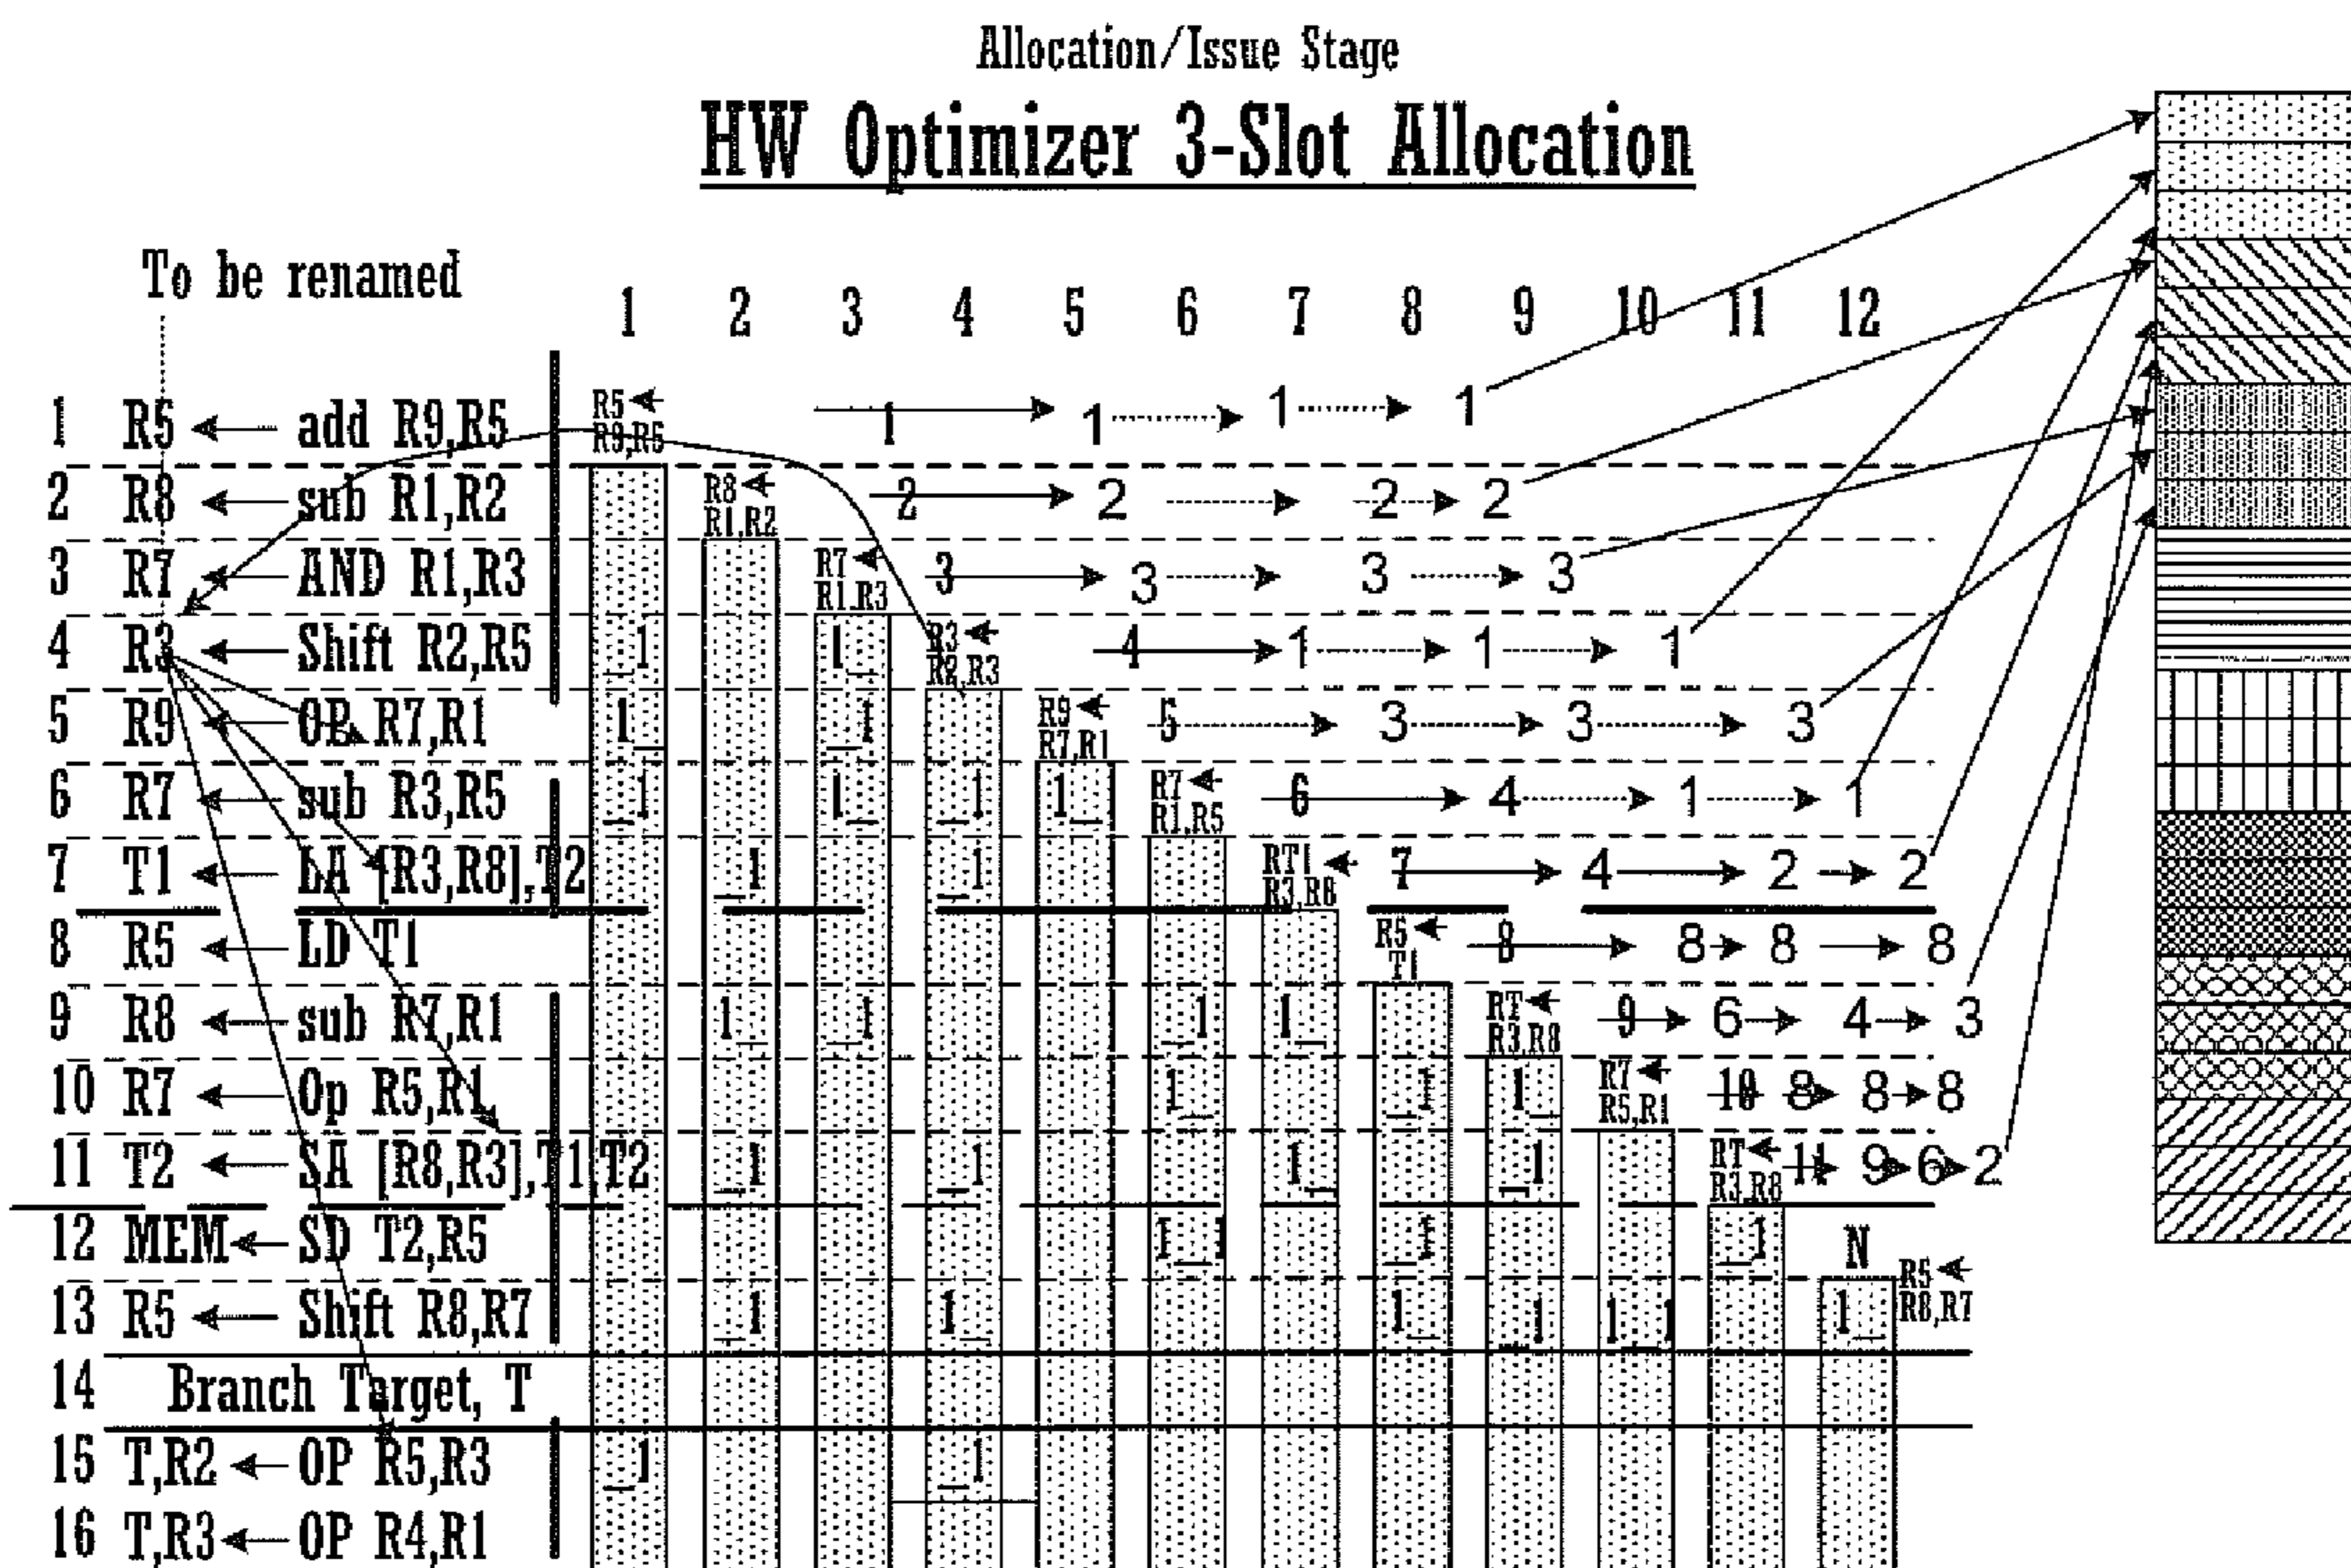
(51) **Int. Cl.**

G06F 9/38 (2006.01)

G06F 9/30 (2006.01)

(57) **ABSTRACT**

A method for accelerating code optimization a microprocessor. The method includes fetching an incoming microinstruction sequence using an instruction fetch component and transferring the fetched macroinstructions to a decoding component for decoding into microinstructions. Optimization processing is performed by reordering the microinstruction sequence into an optimized microinstruction sequence comprising a plurality of dependent code groups. The optimized microinstruction sequence is output to a microprocessor pipeline for execution. A copy of the optimized microinstruction sequence is stored into a sequence cache for subsequent use upon a subsequent hit optimized microinstruction sequence.



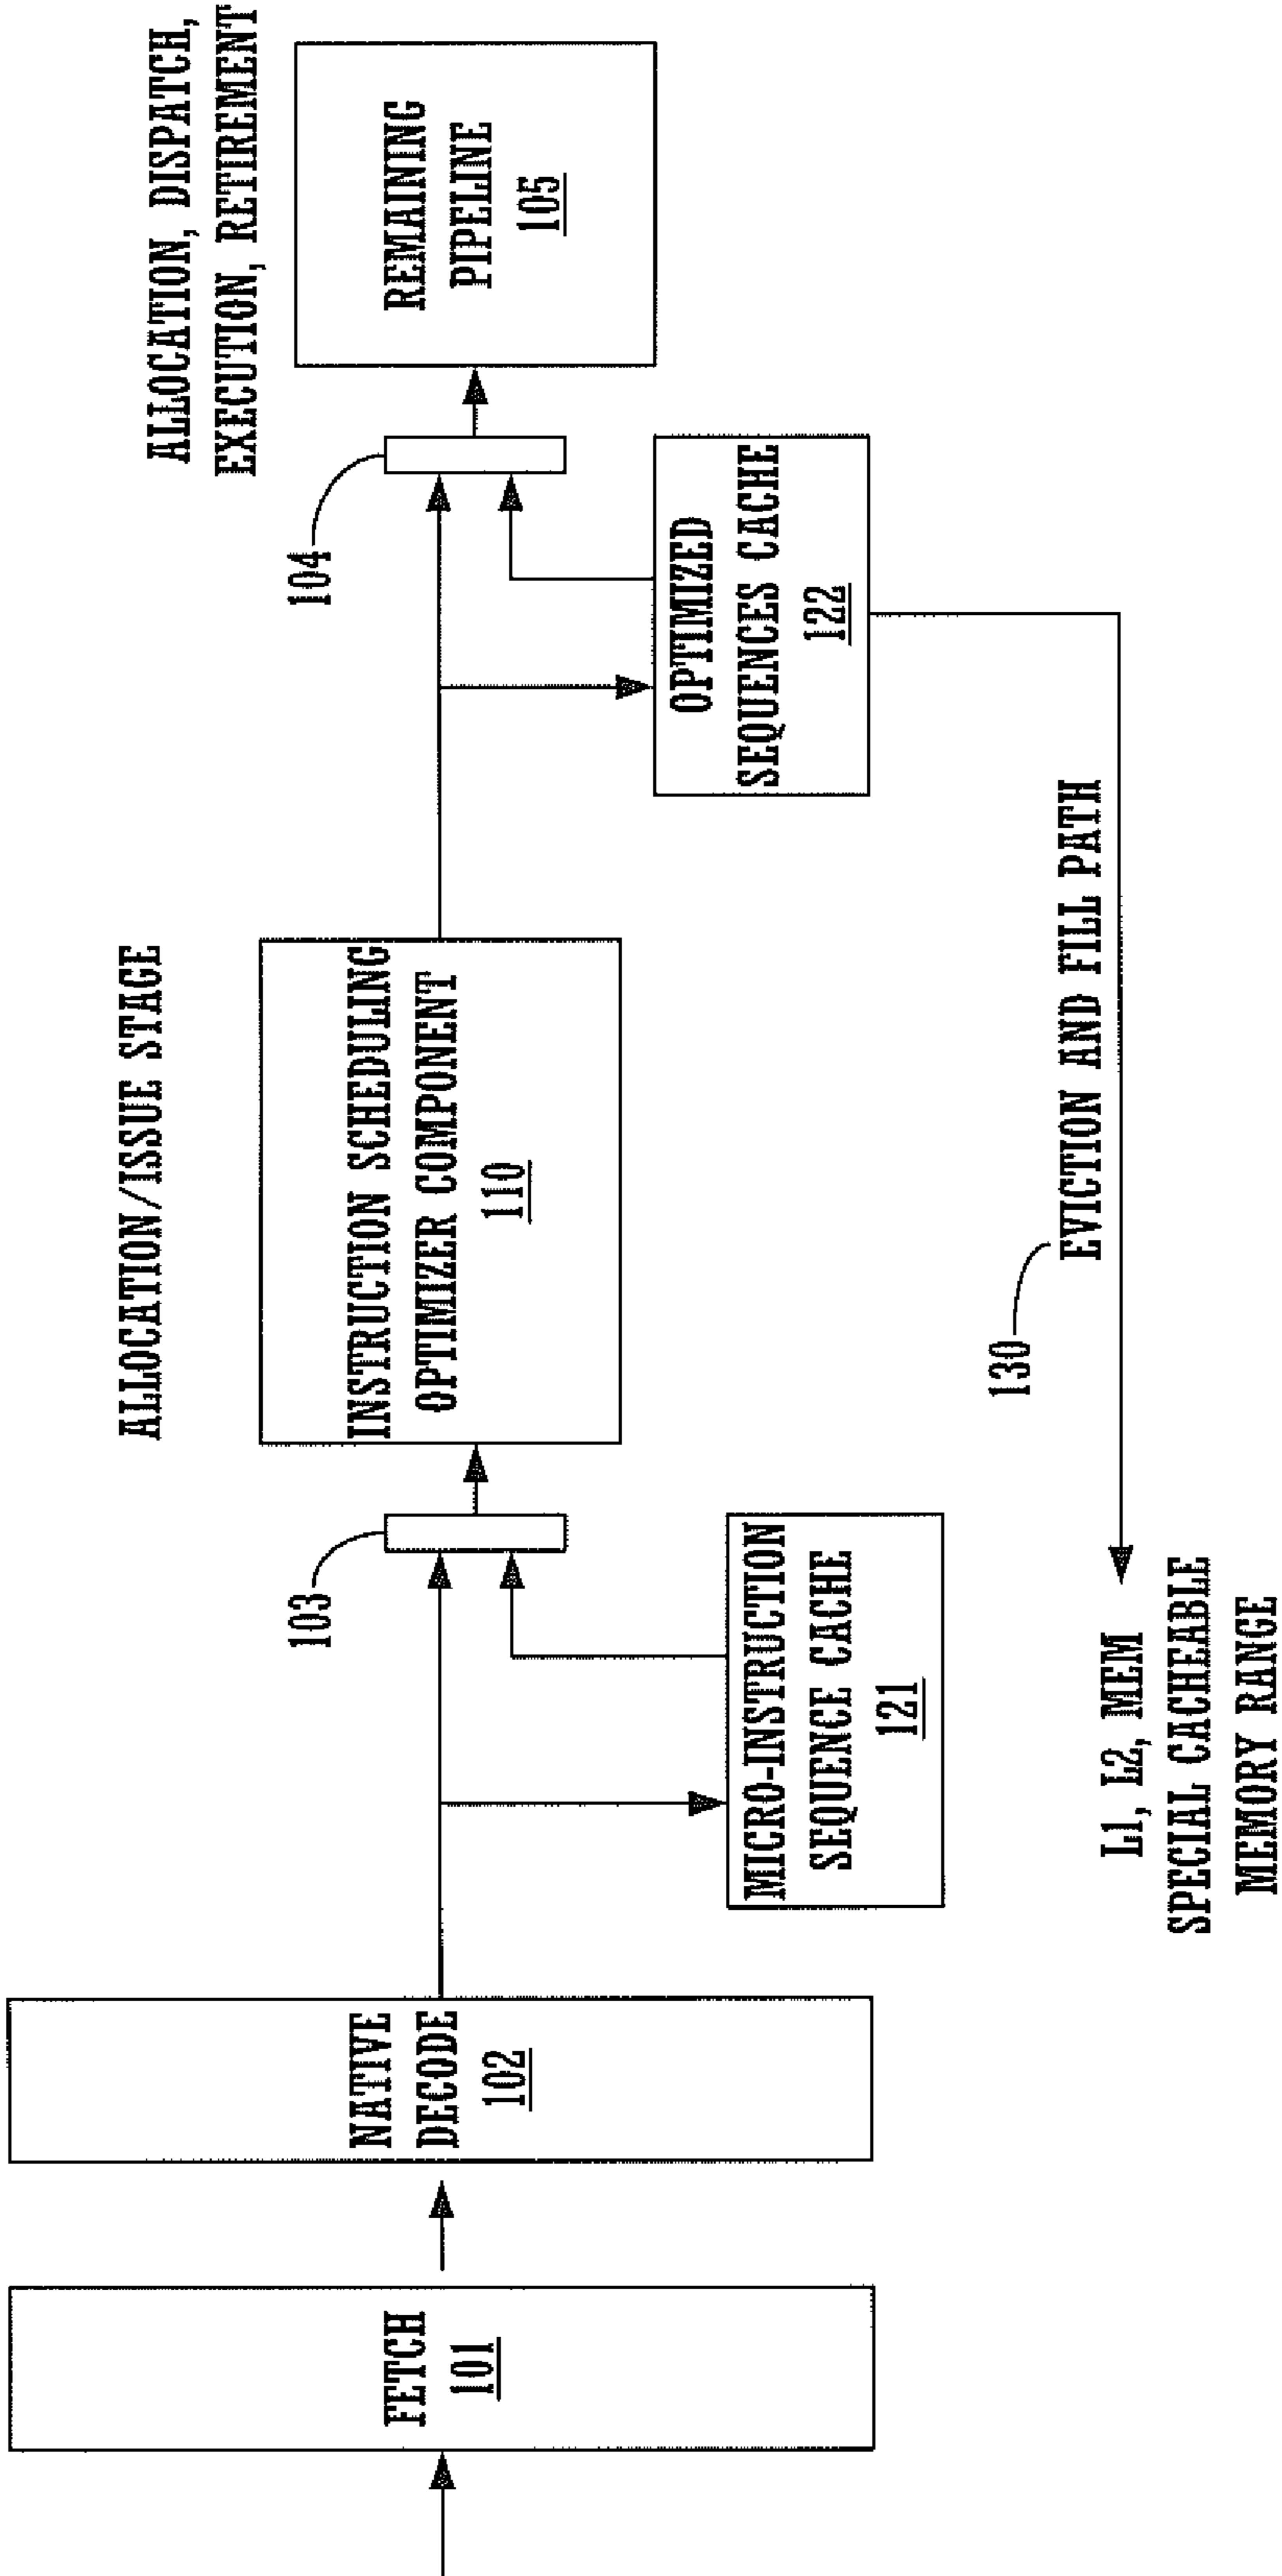
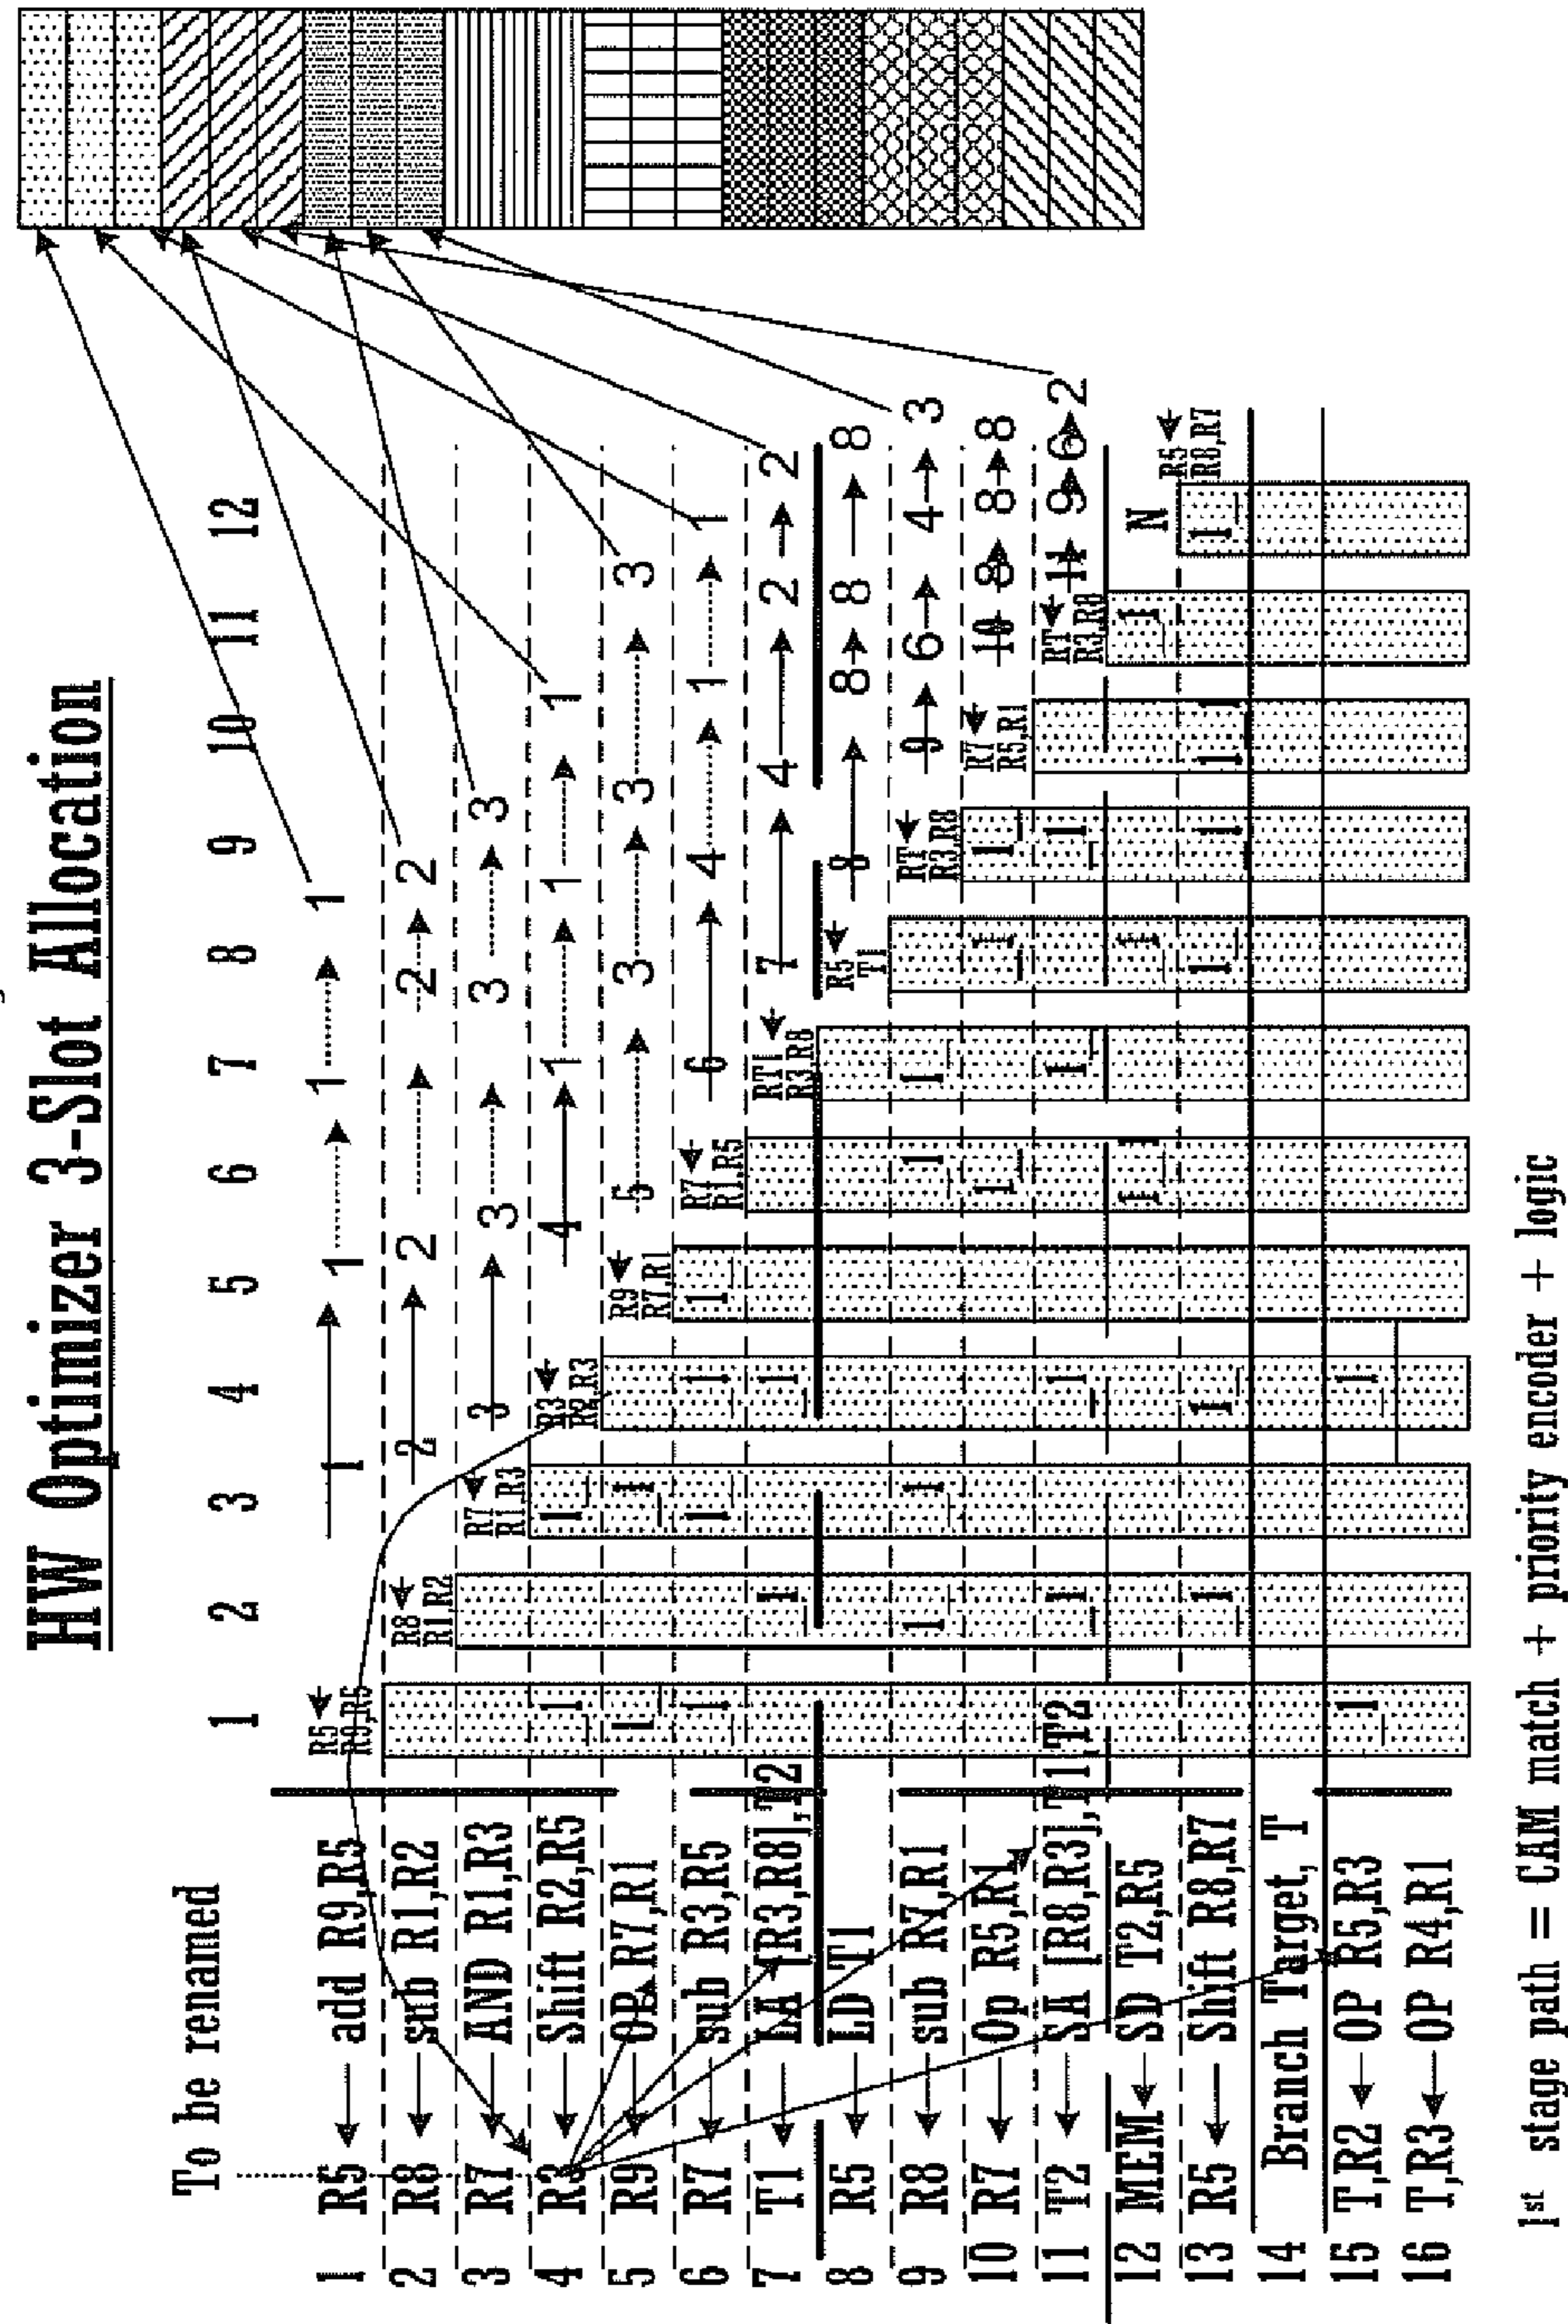


FIGURE 1

Allocation/Issue Stage

HW Optimizer 3-Slot Allocation

HW Optimizer 3-Slot Allocation



1st stage path = CAM match + priority encoder + logic

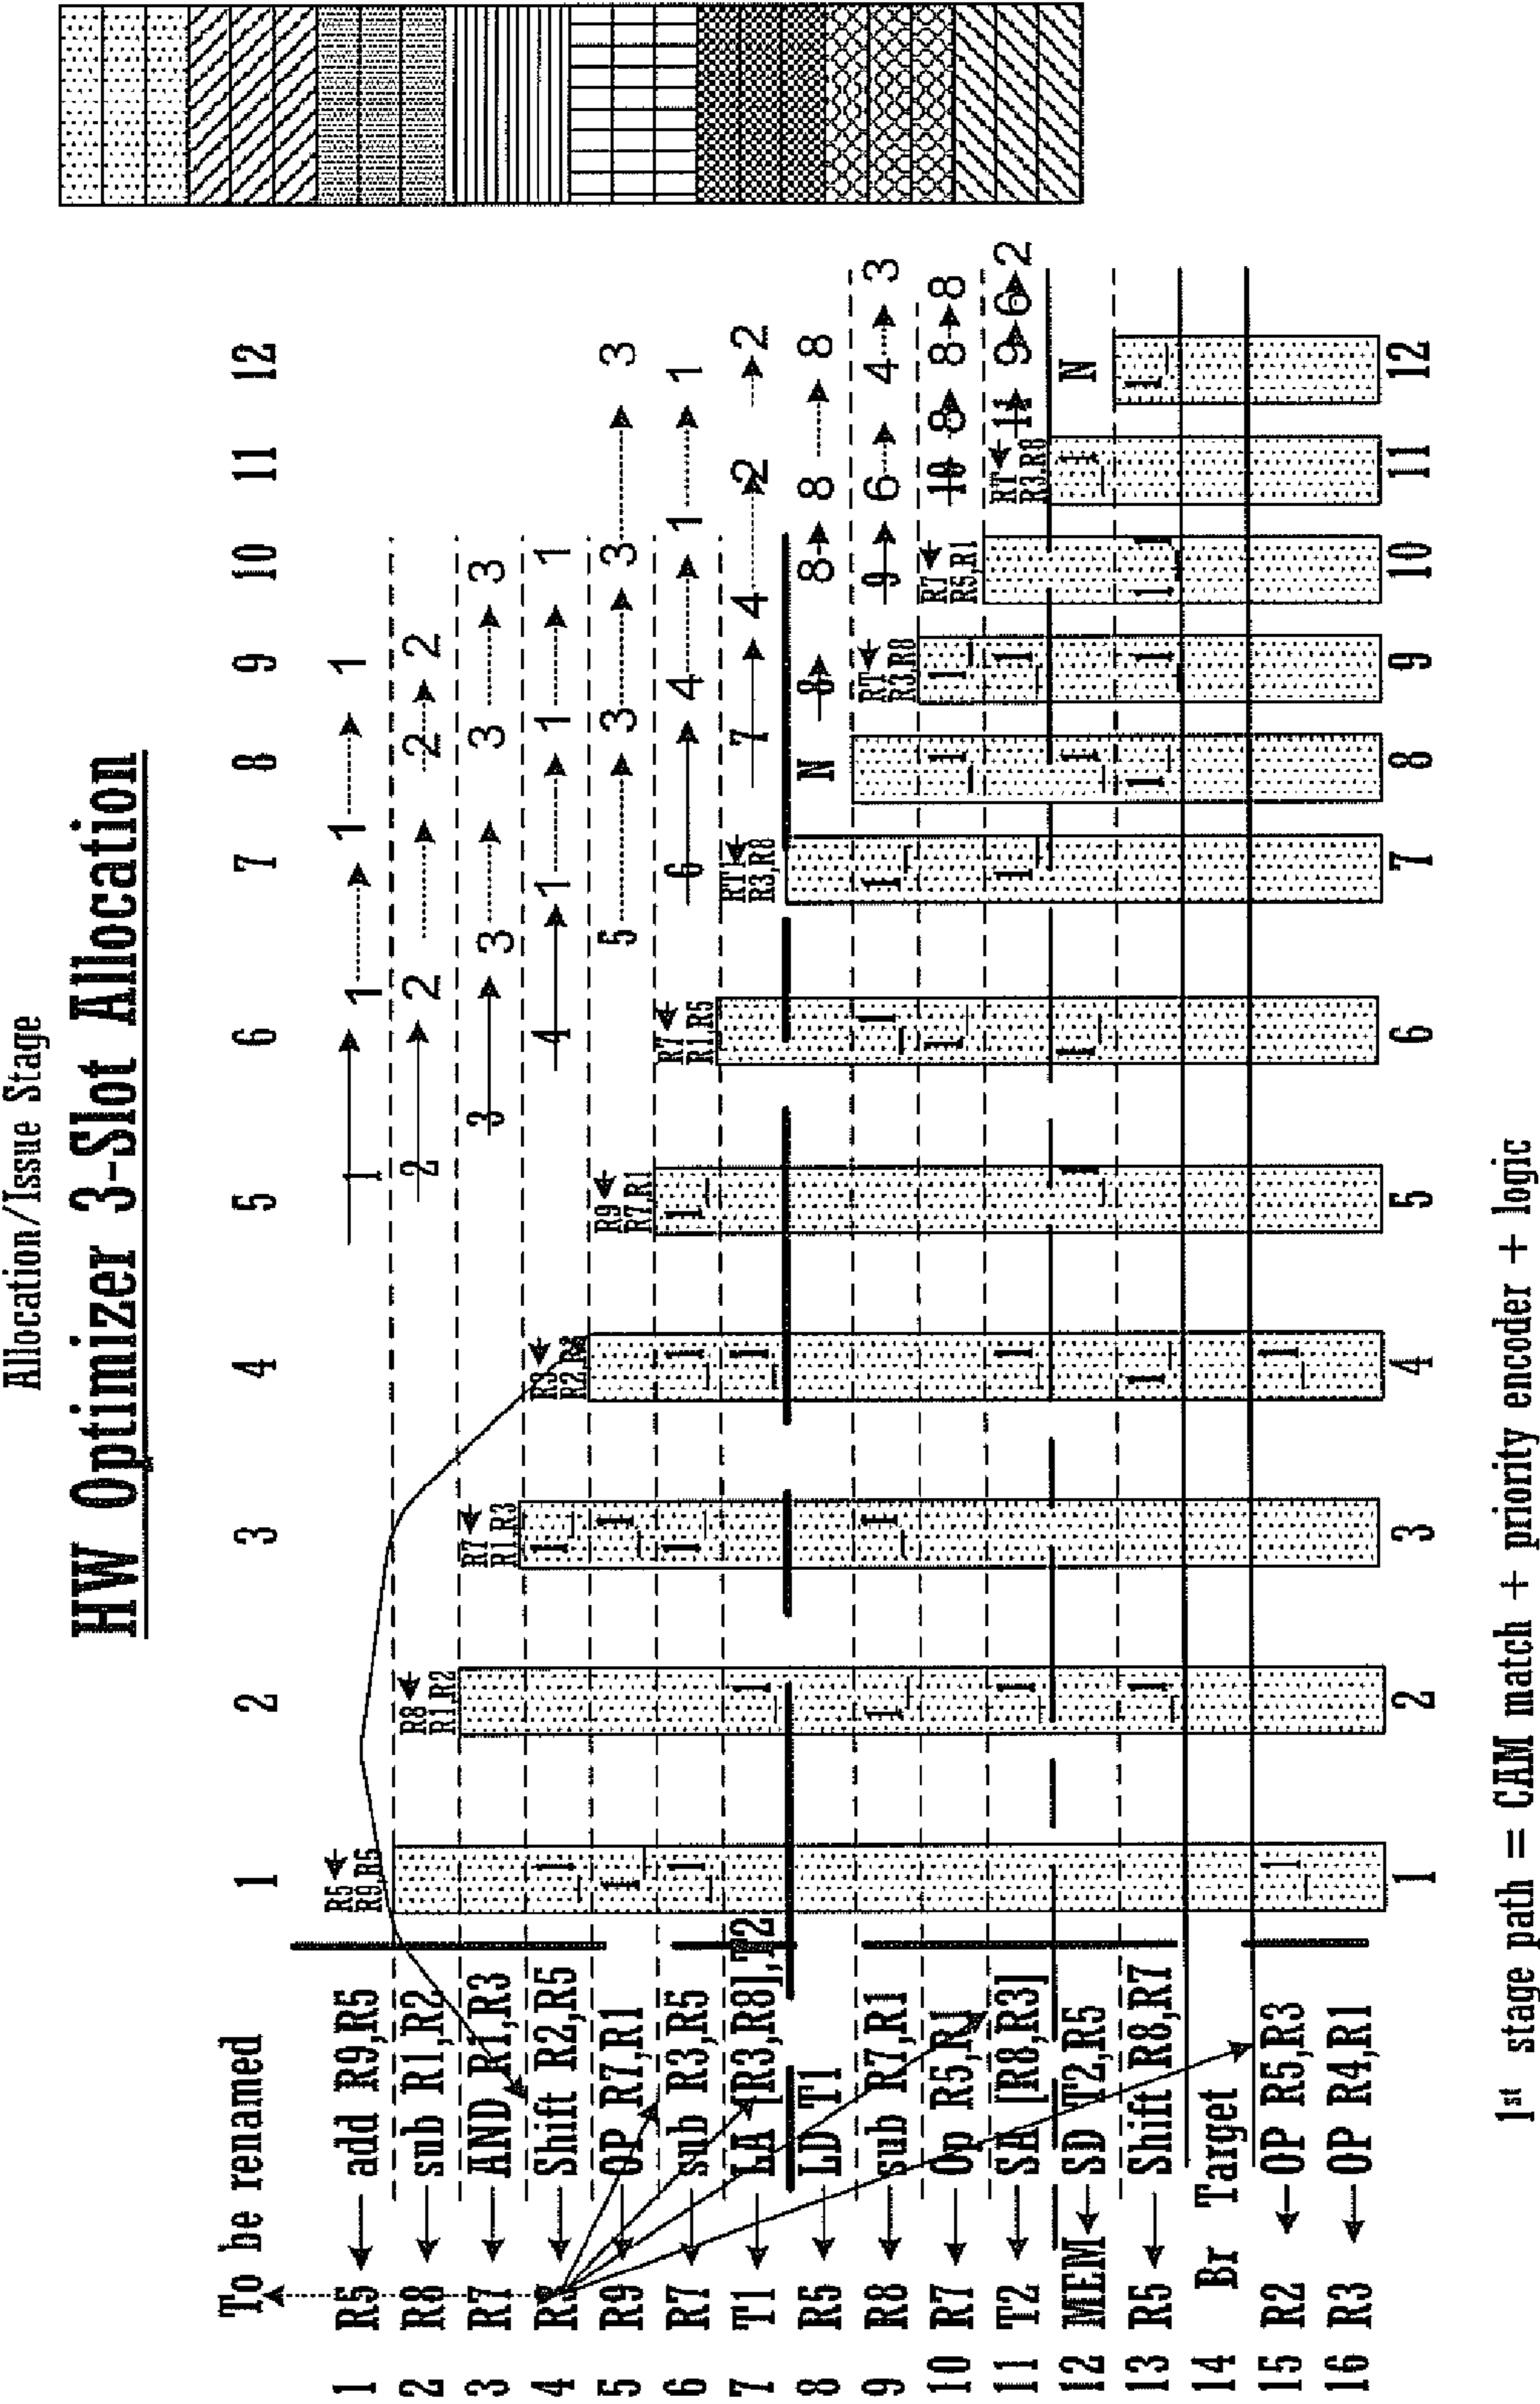
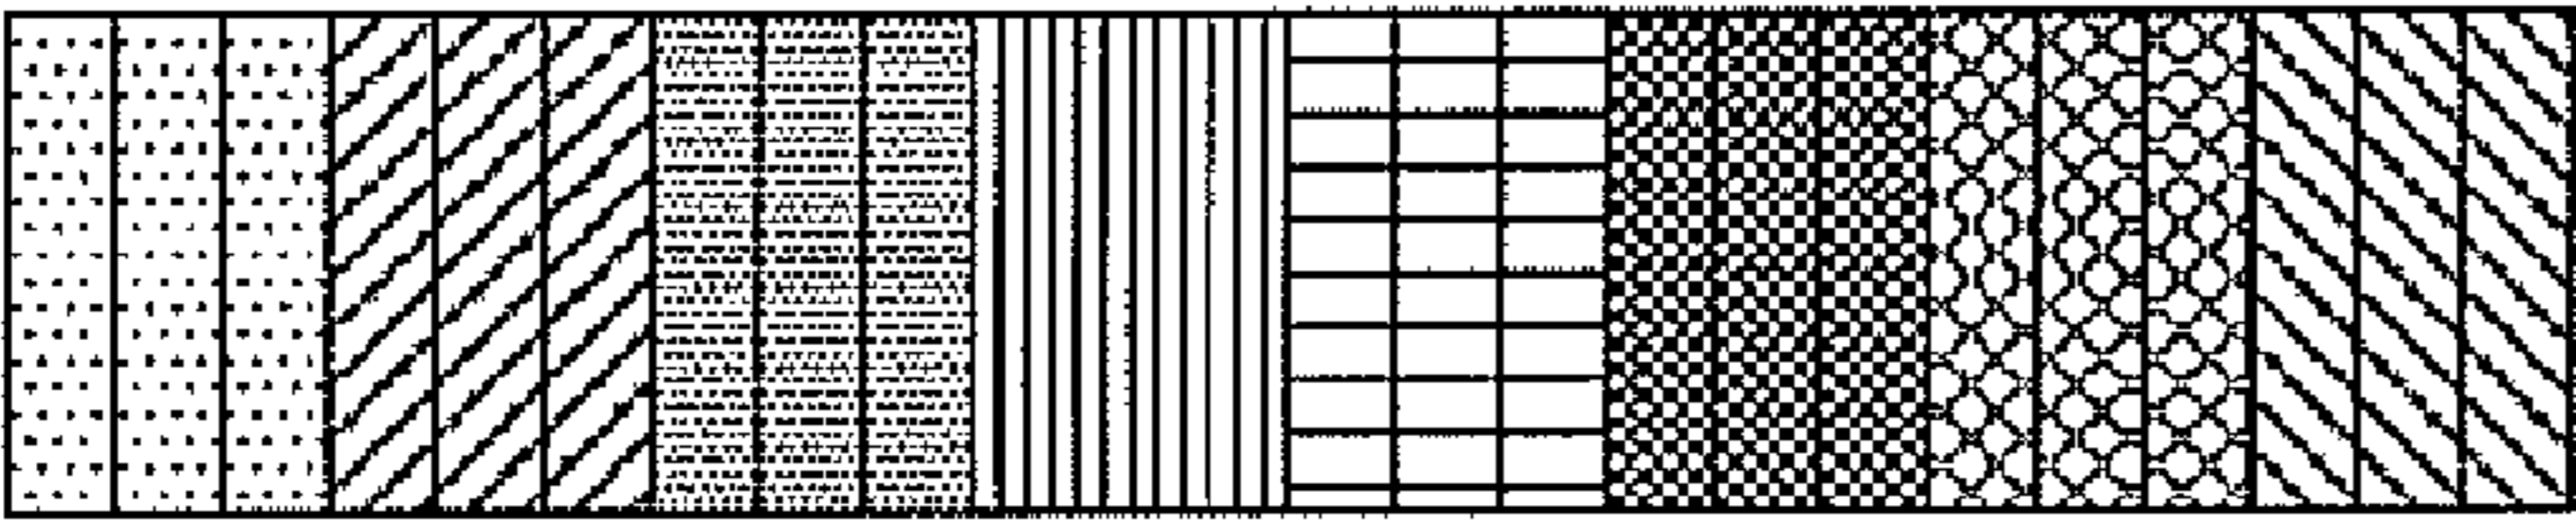


FIGURE 3



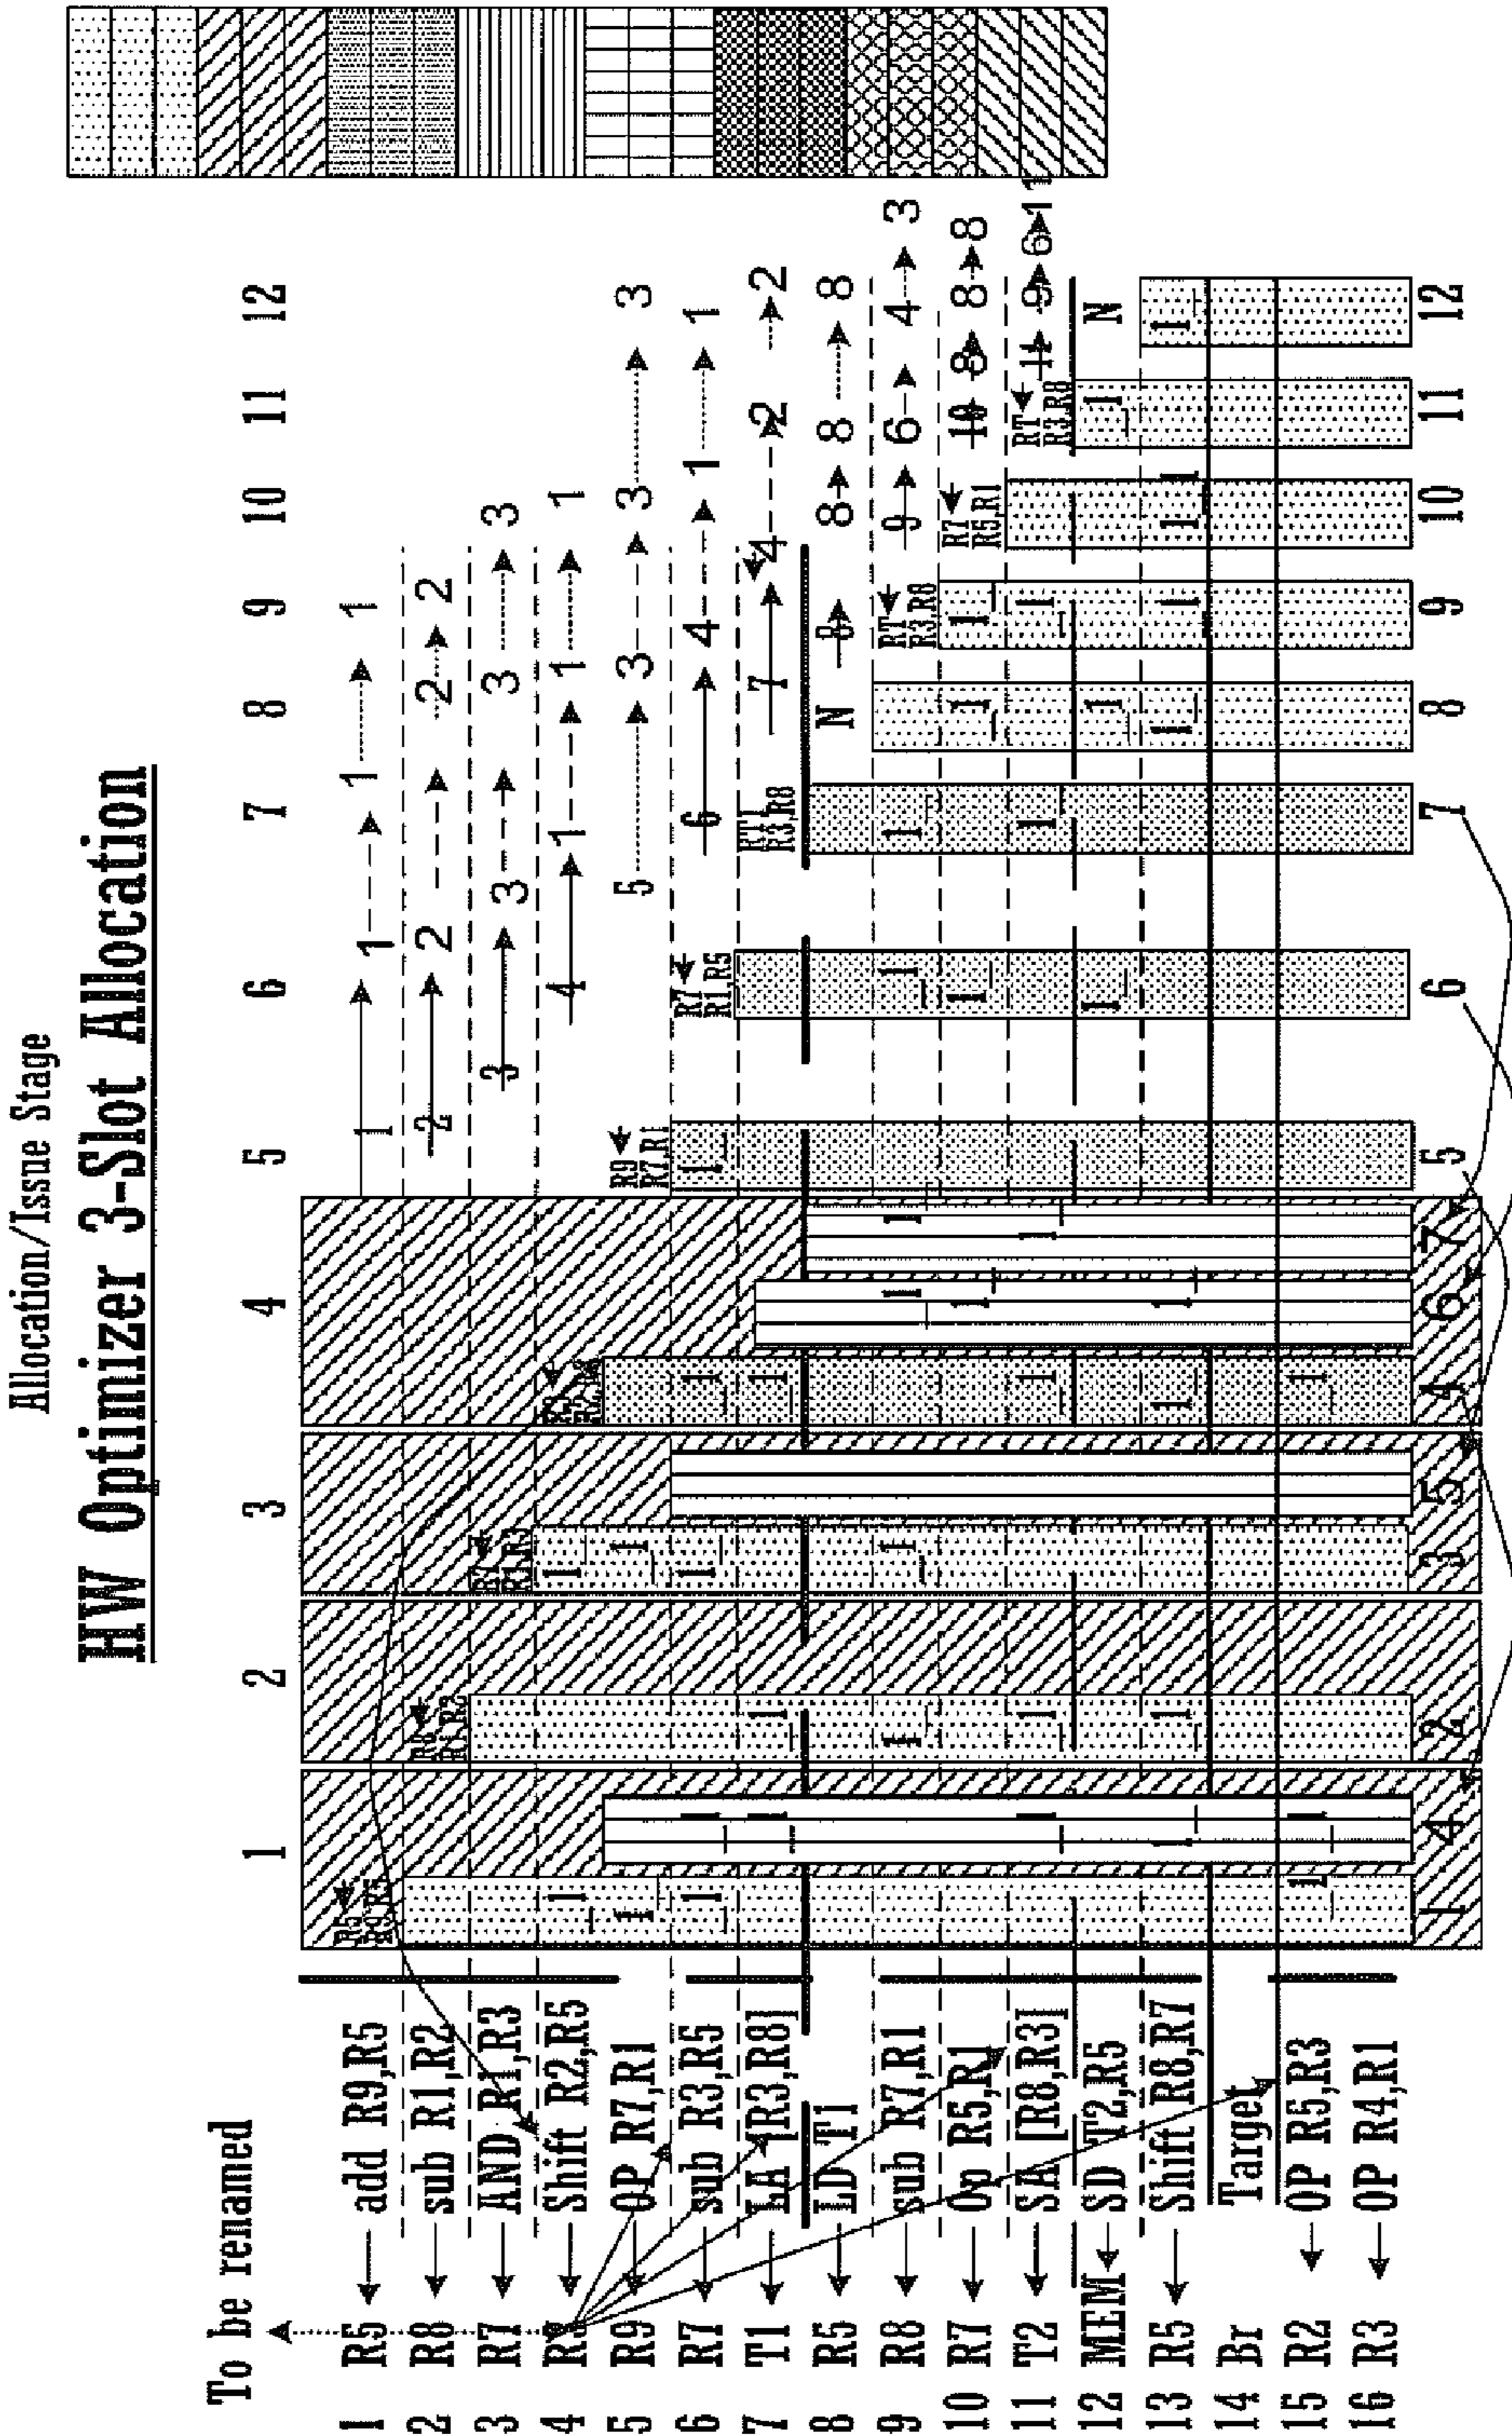
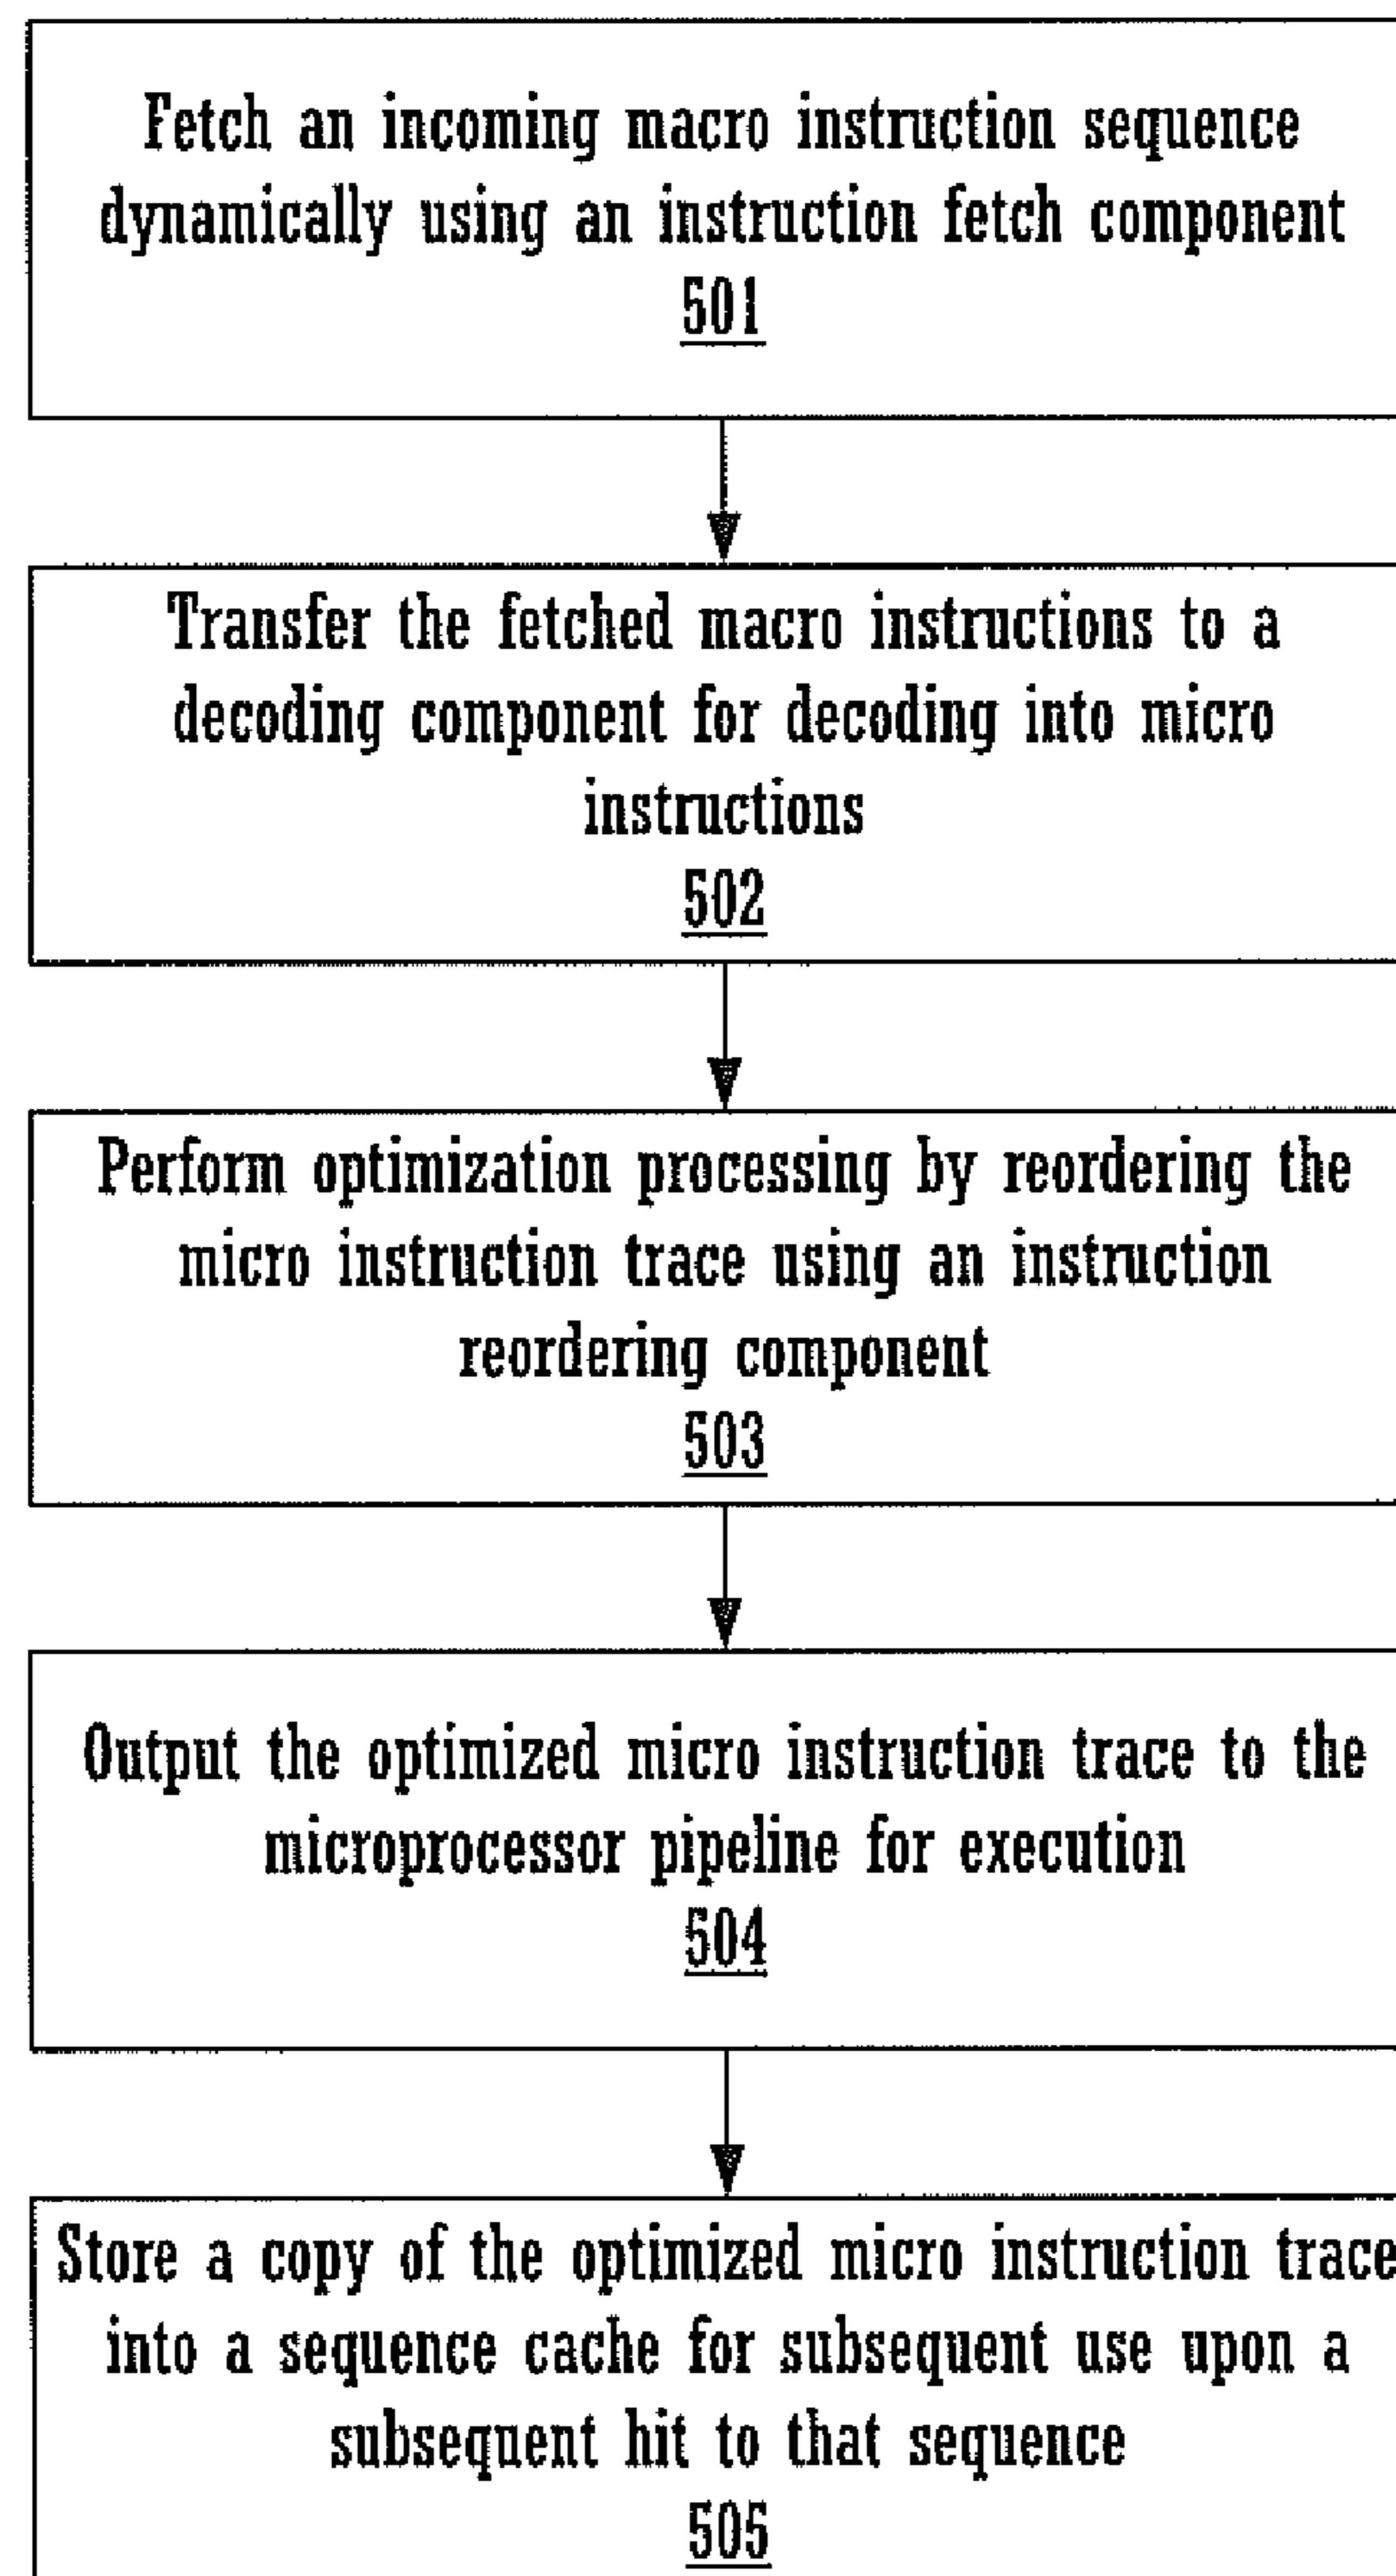
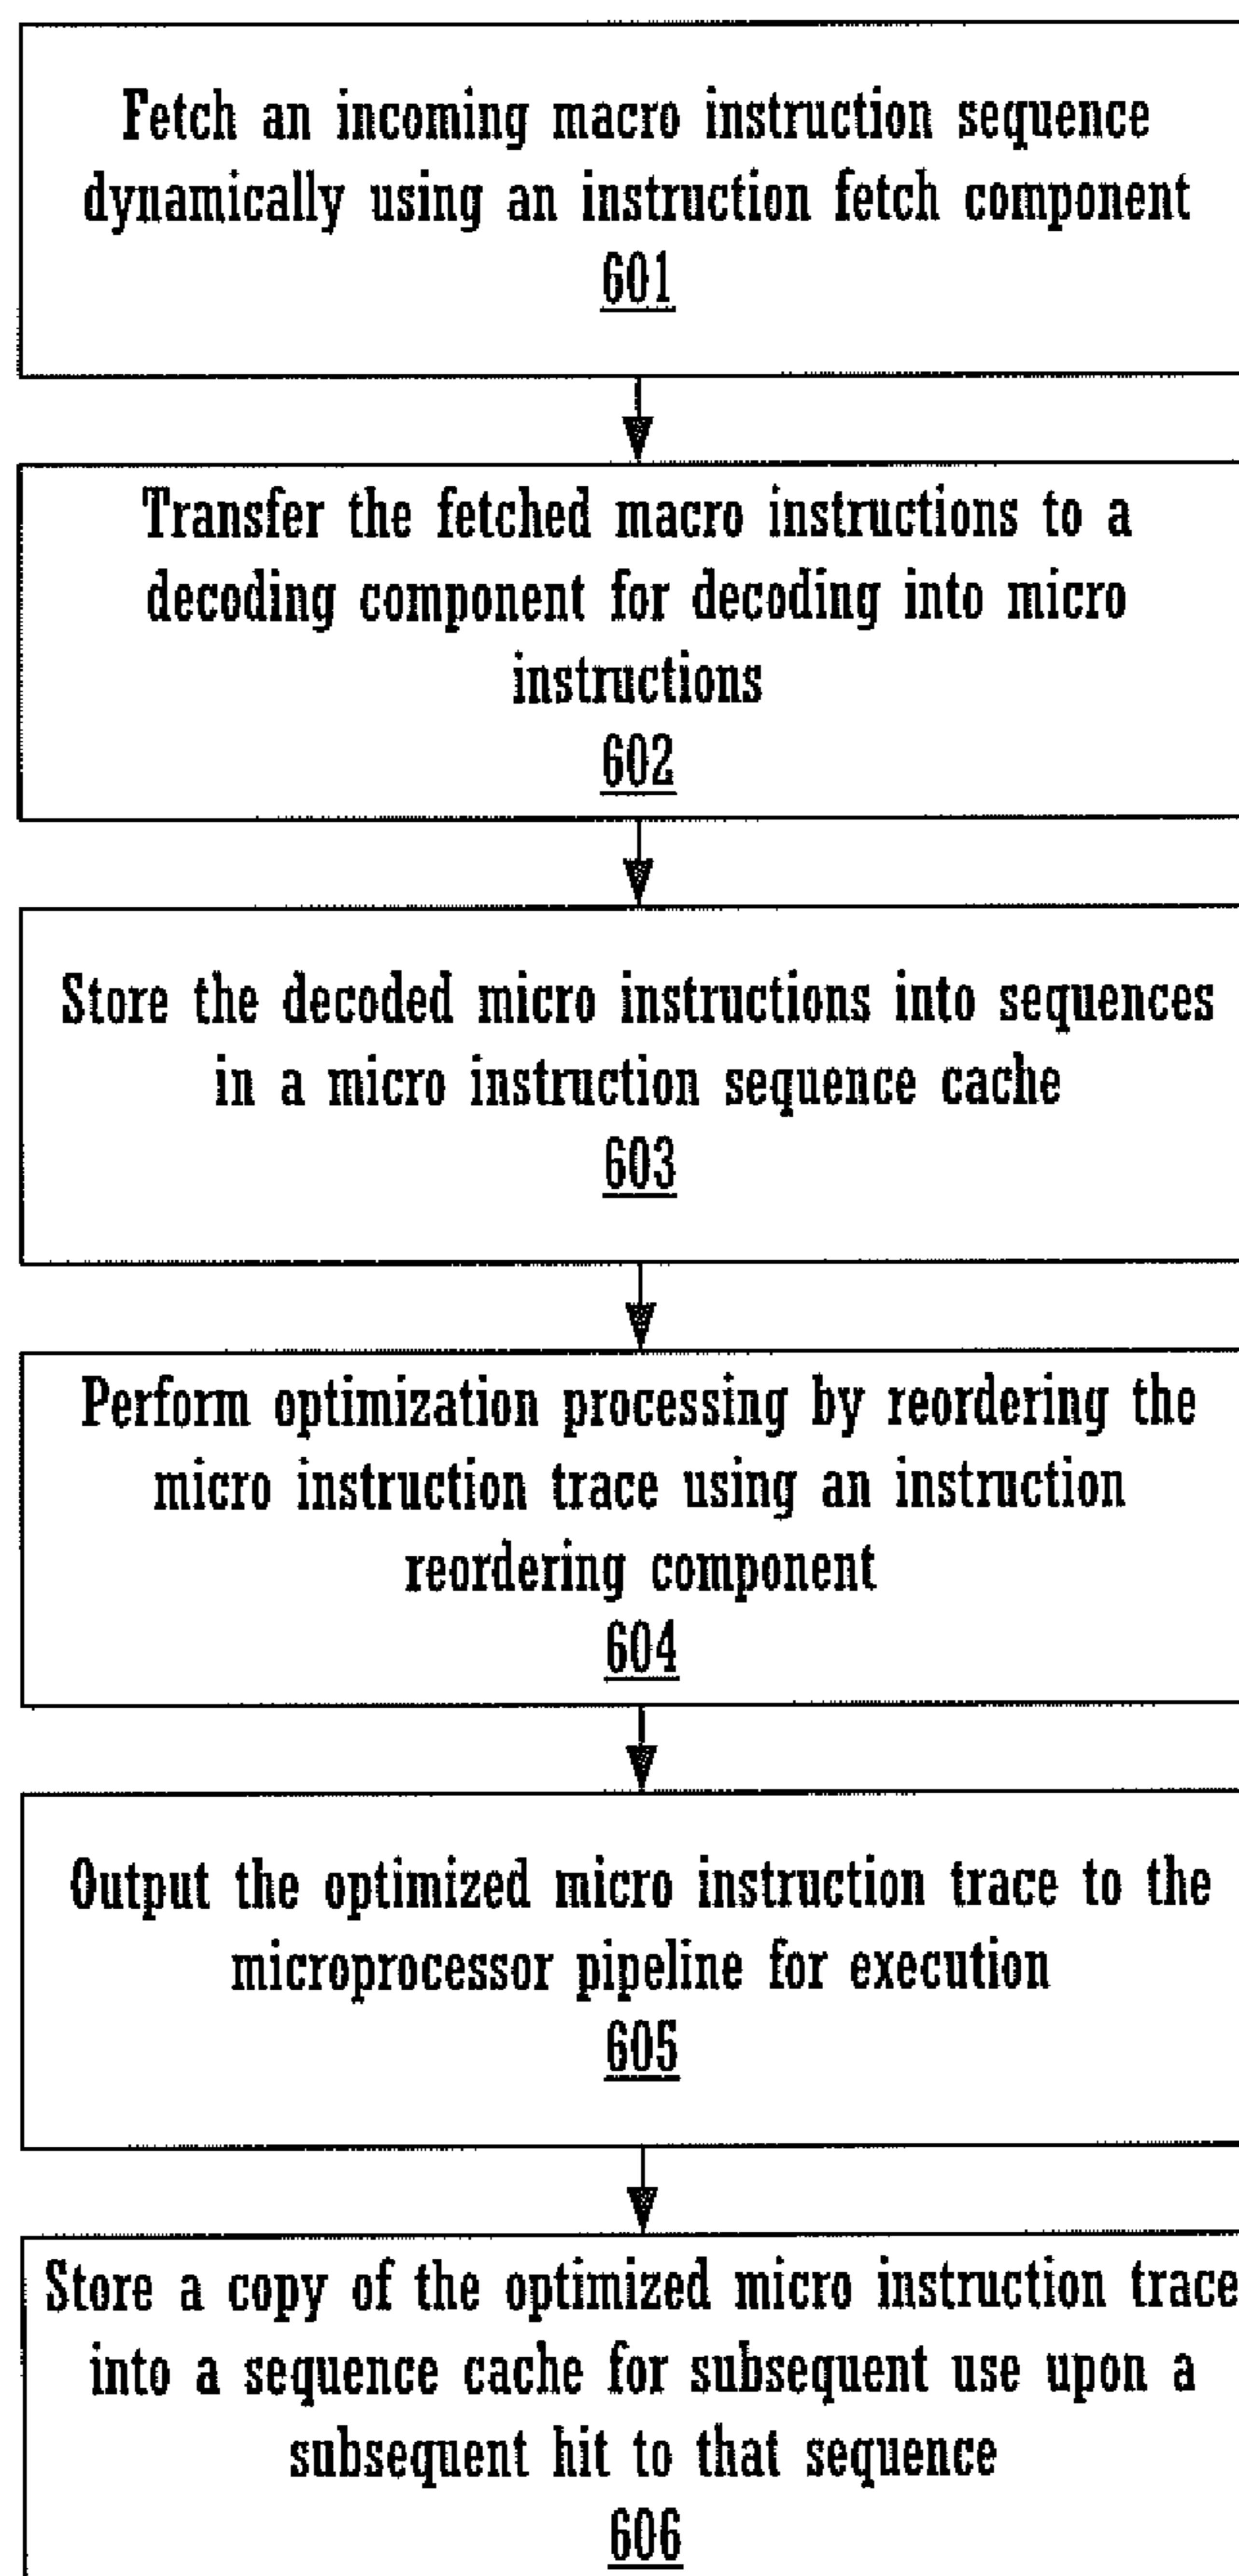


FIGURE 4

500**Hardware Optimization Flowchart****FIGURE 5**

600

Hardware Optimization Flowchart

**FIGURE 6**

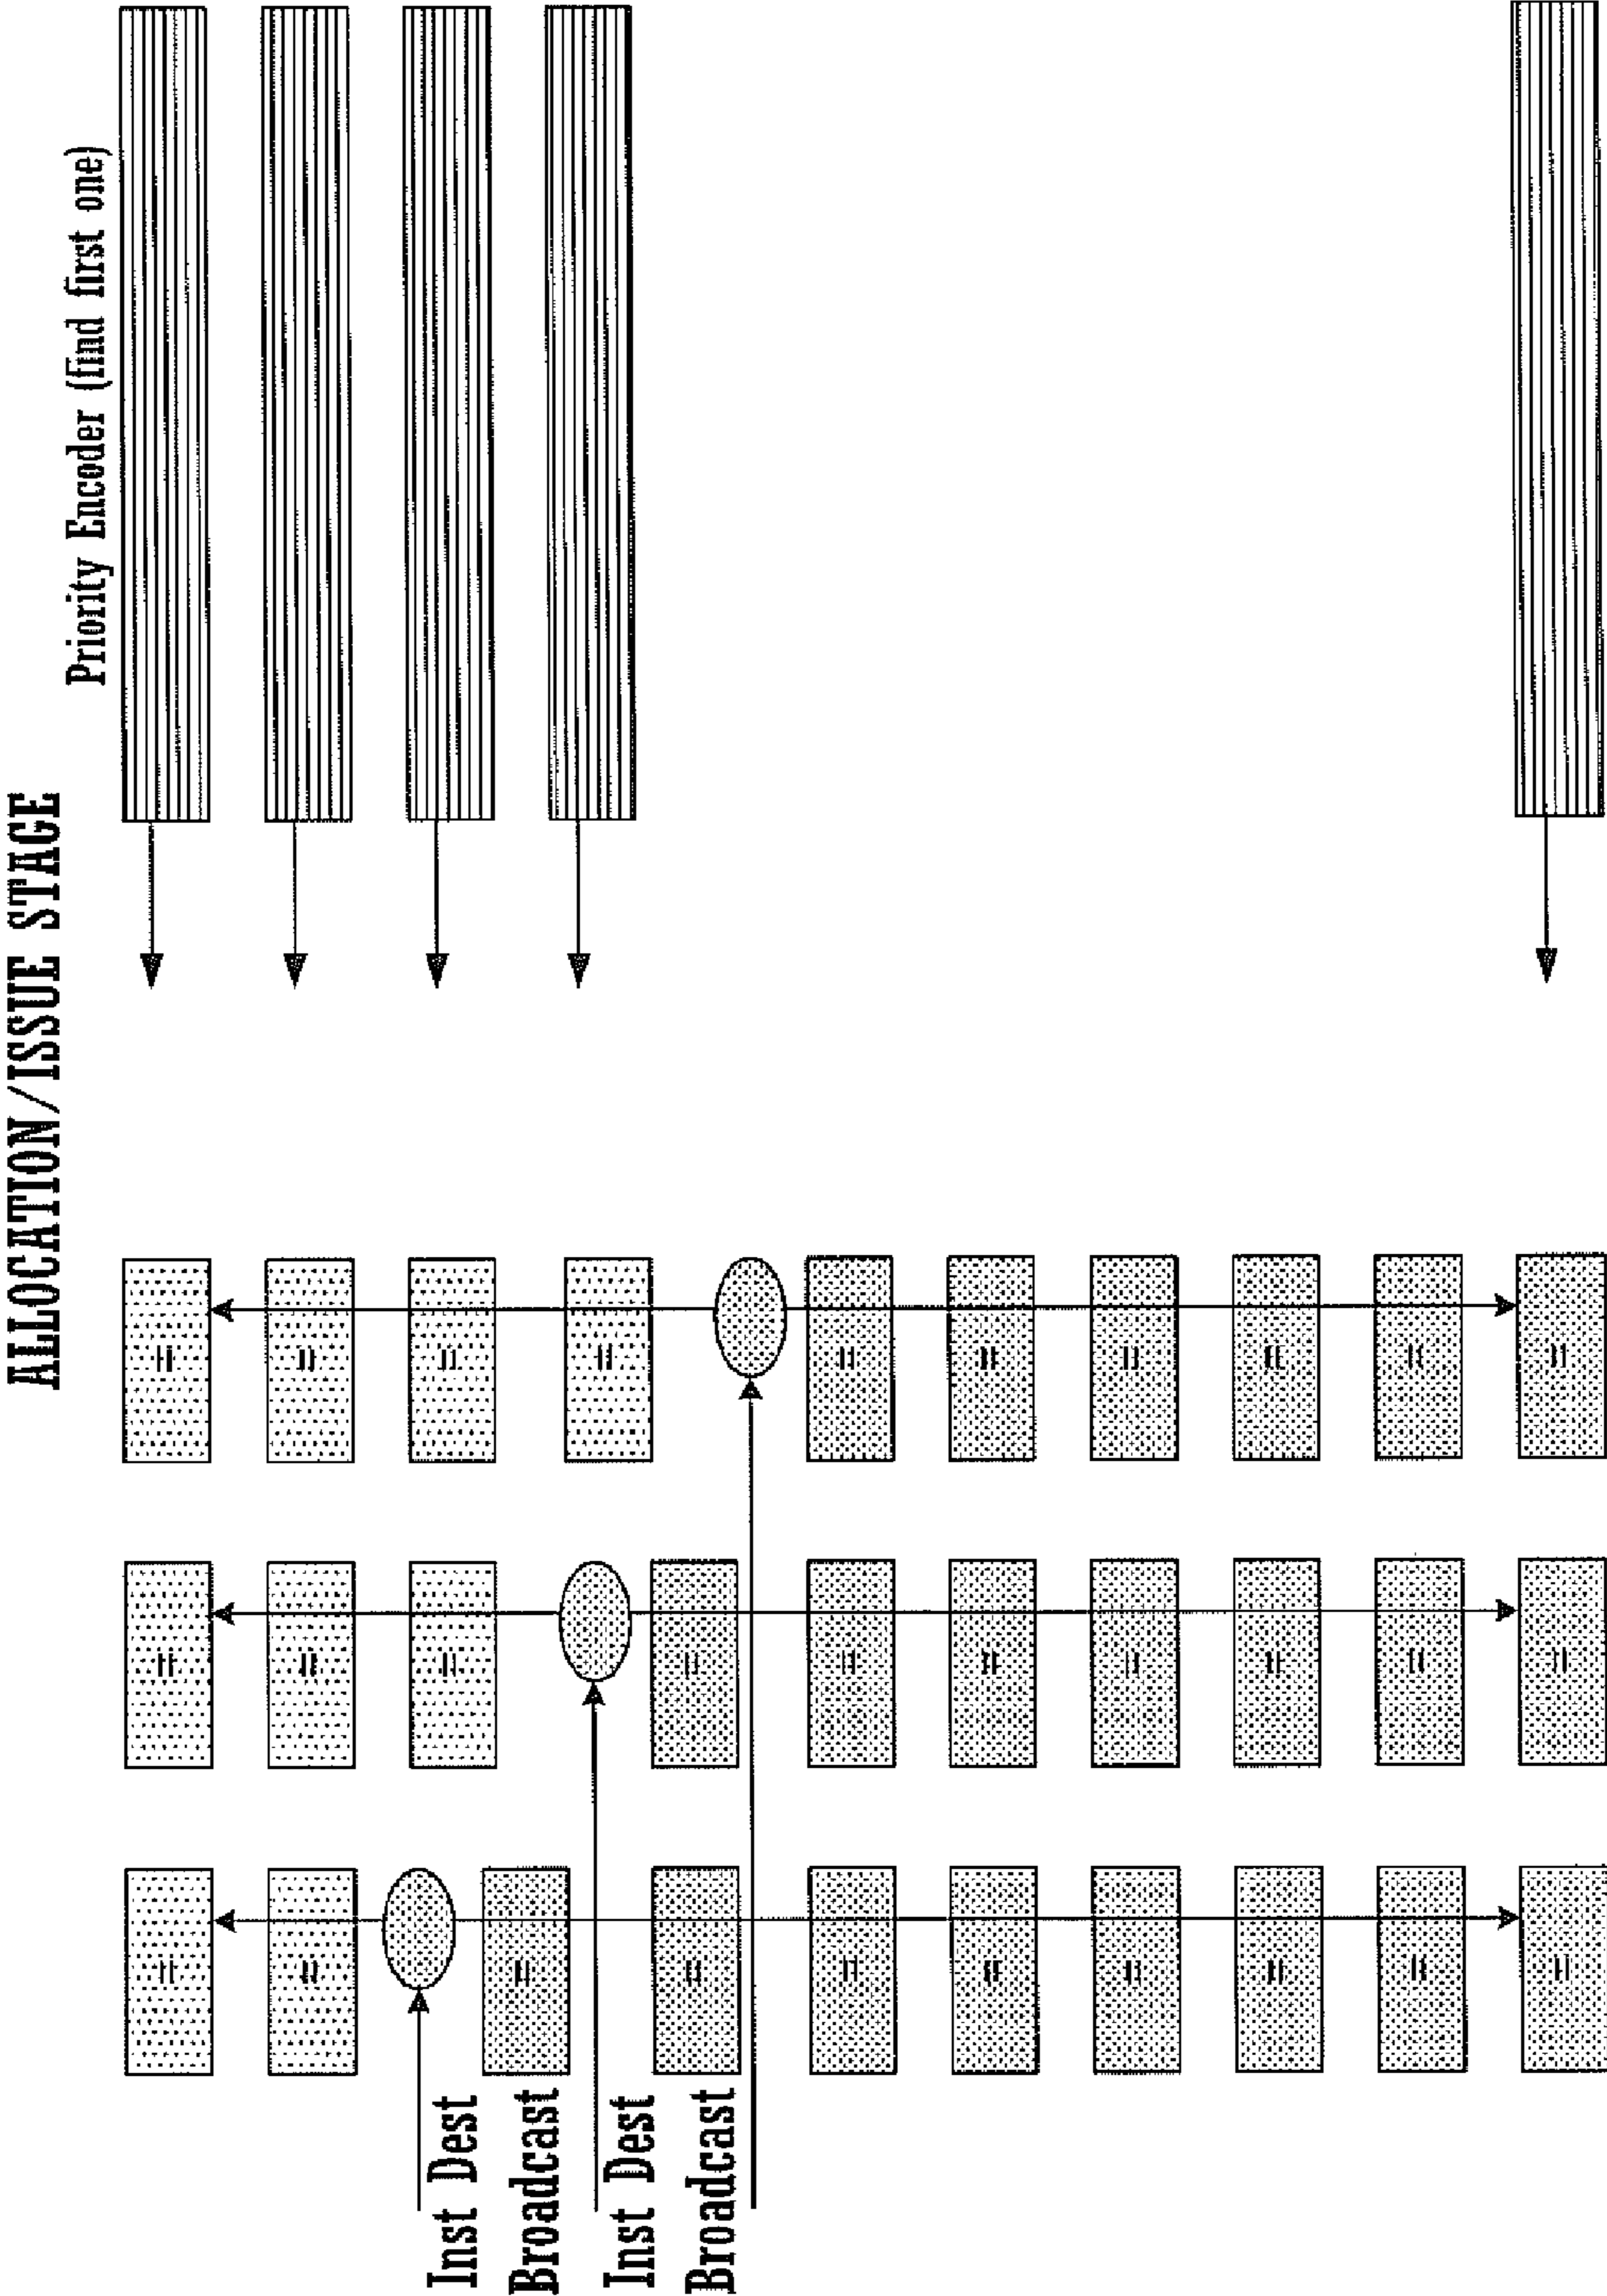
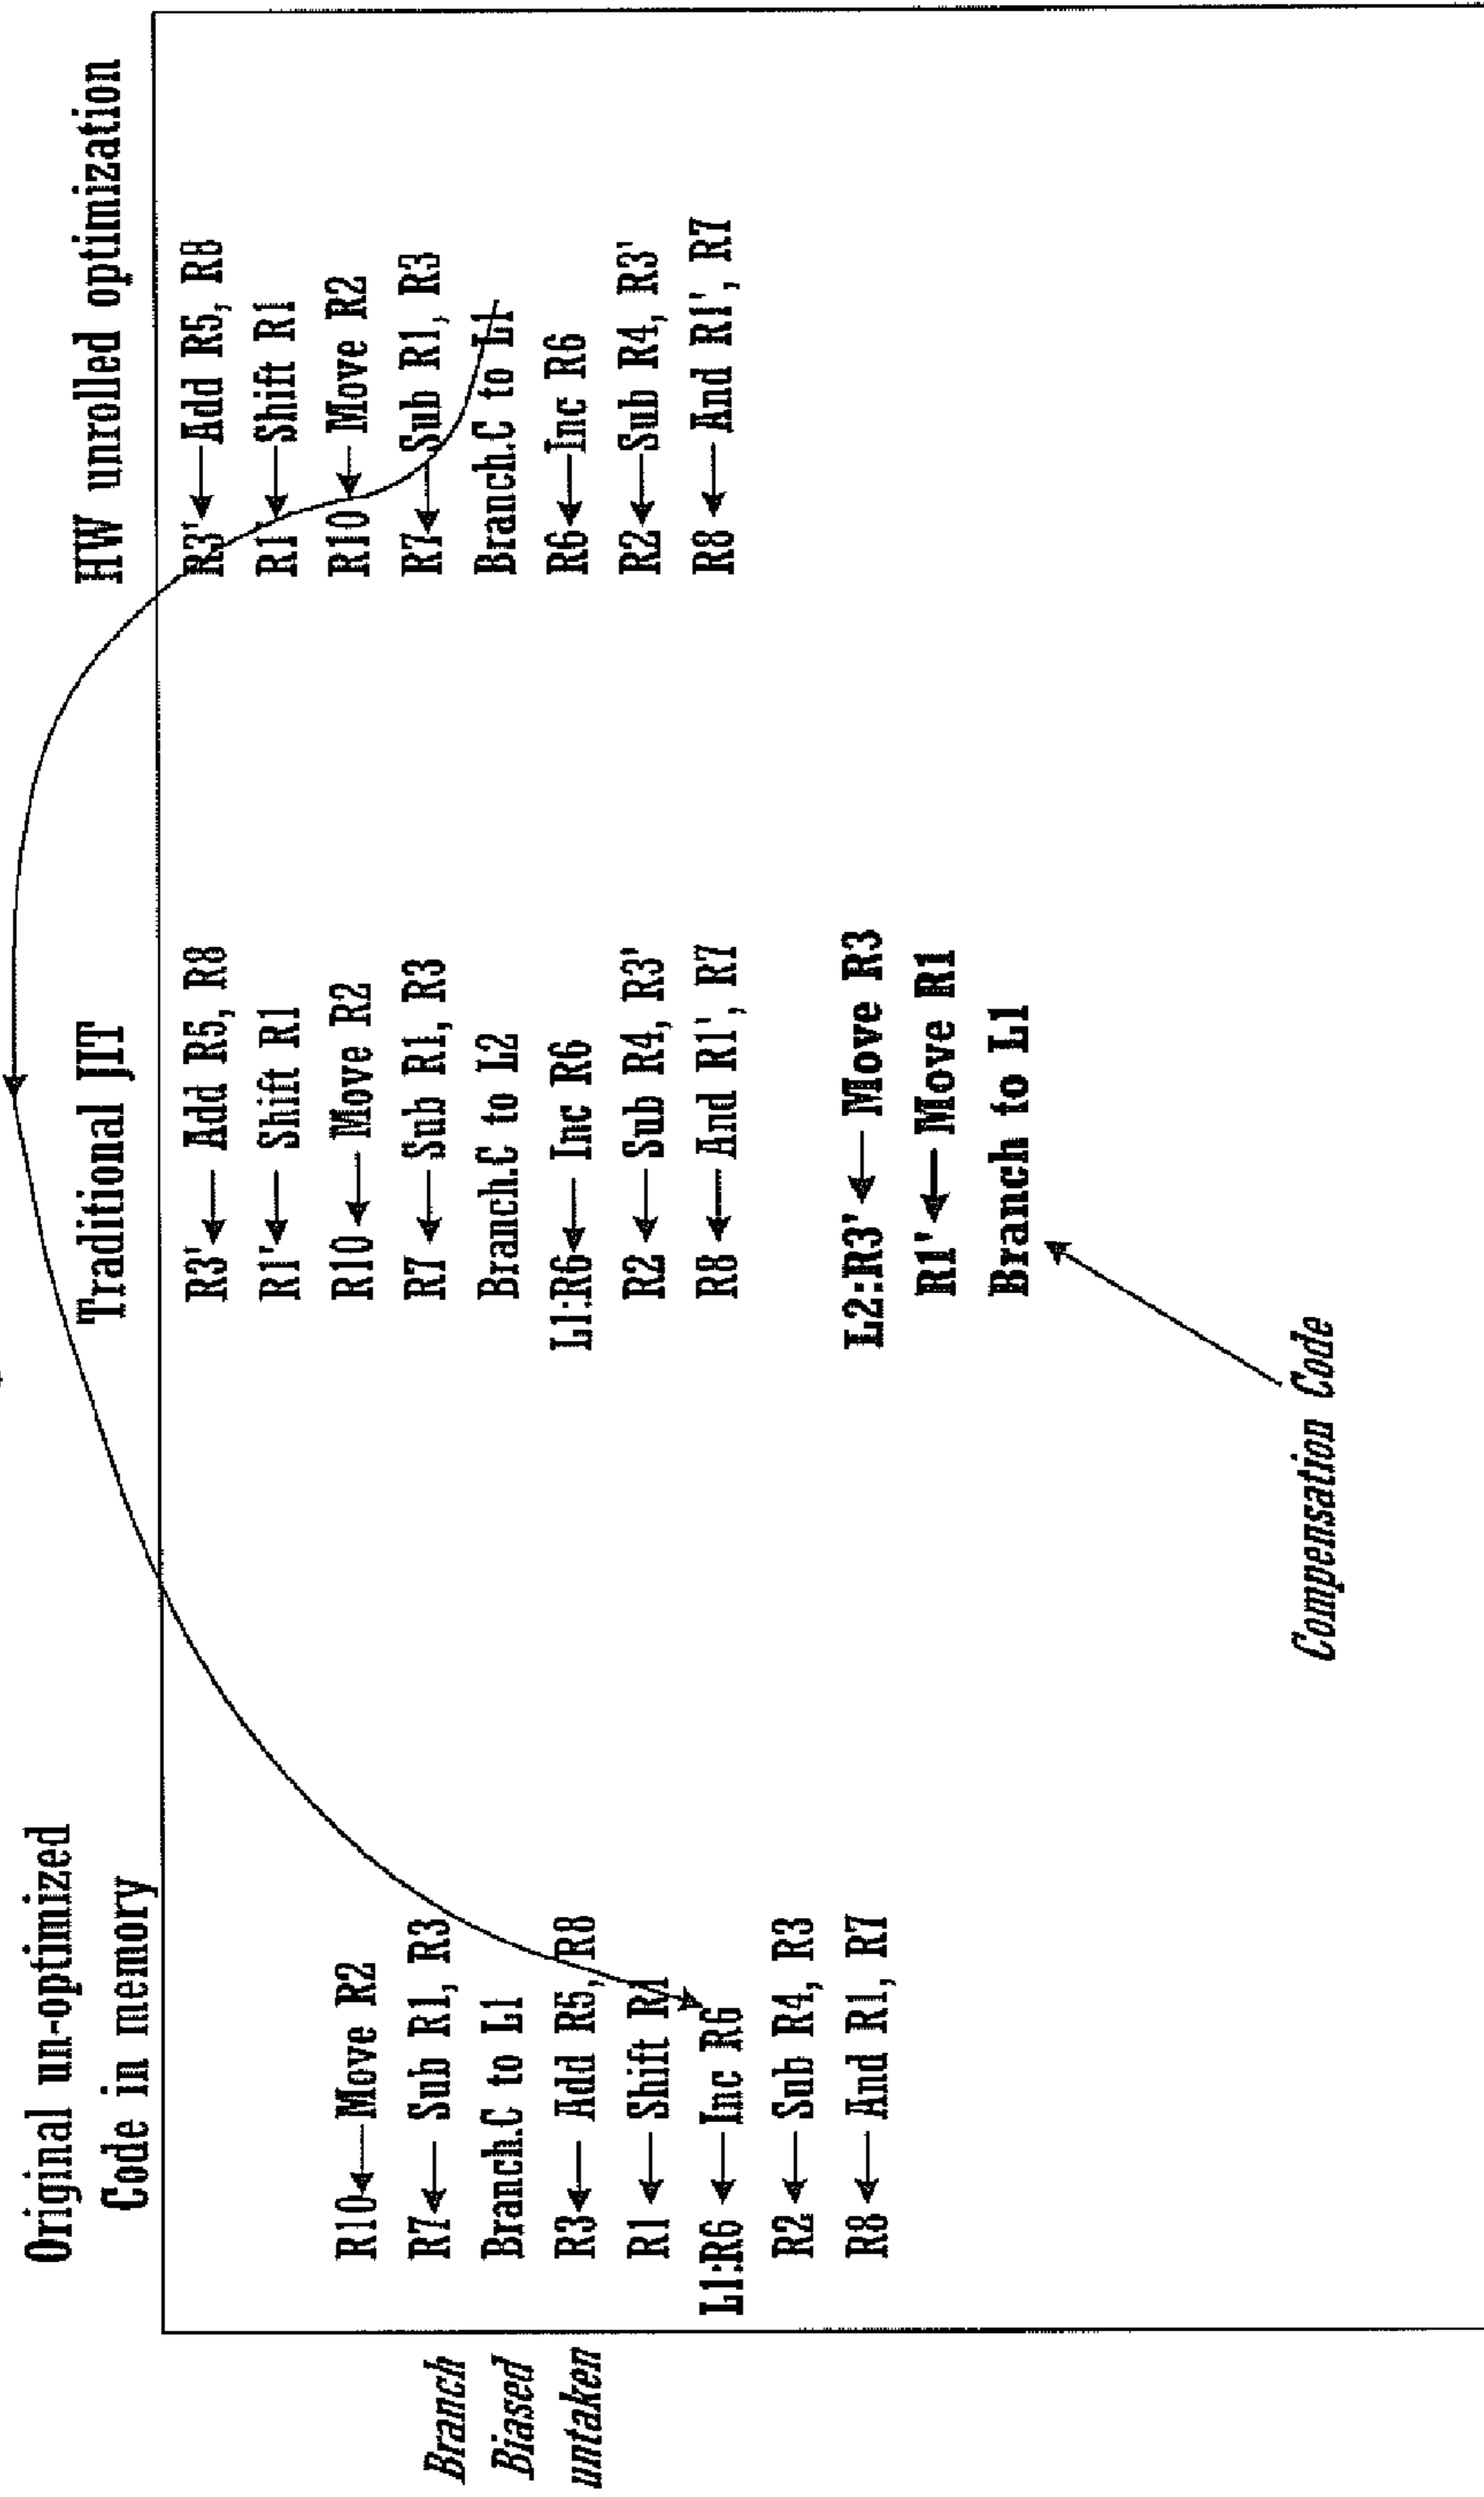


FIGURE 7

HW unrolled Optimization V.S. JIT compilers

Scheduling a head of a branch



HW unrolled Optimization V.S. JIT compilers

Scheduling a head of a store

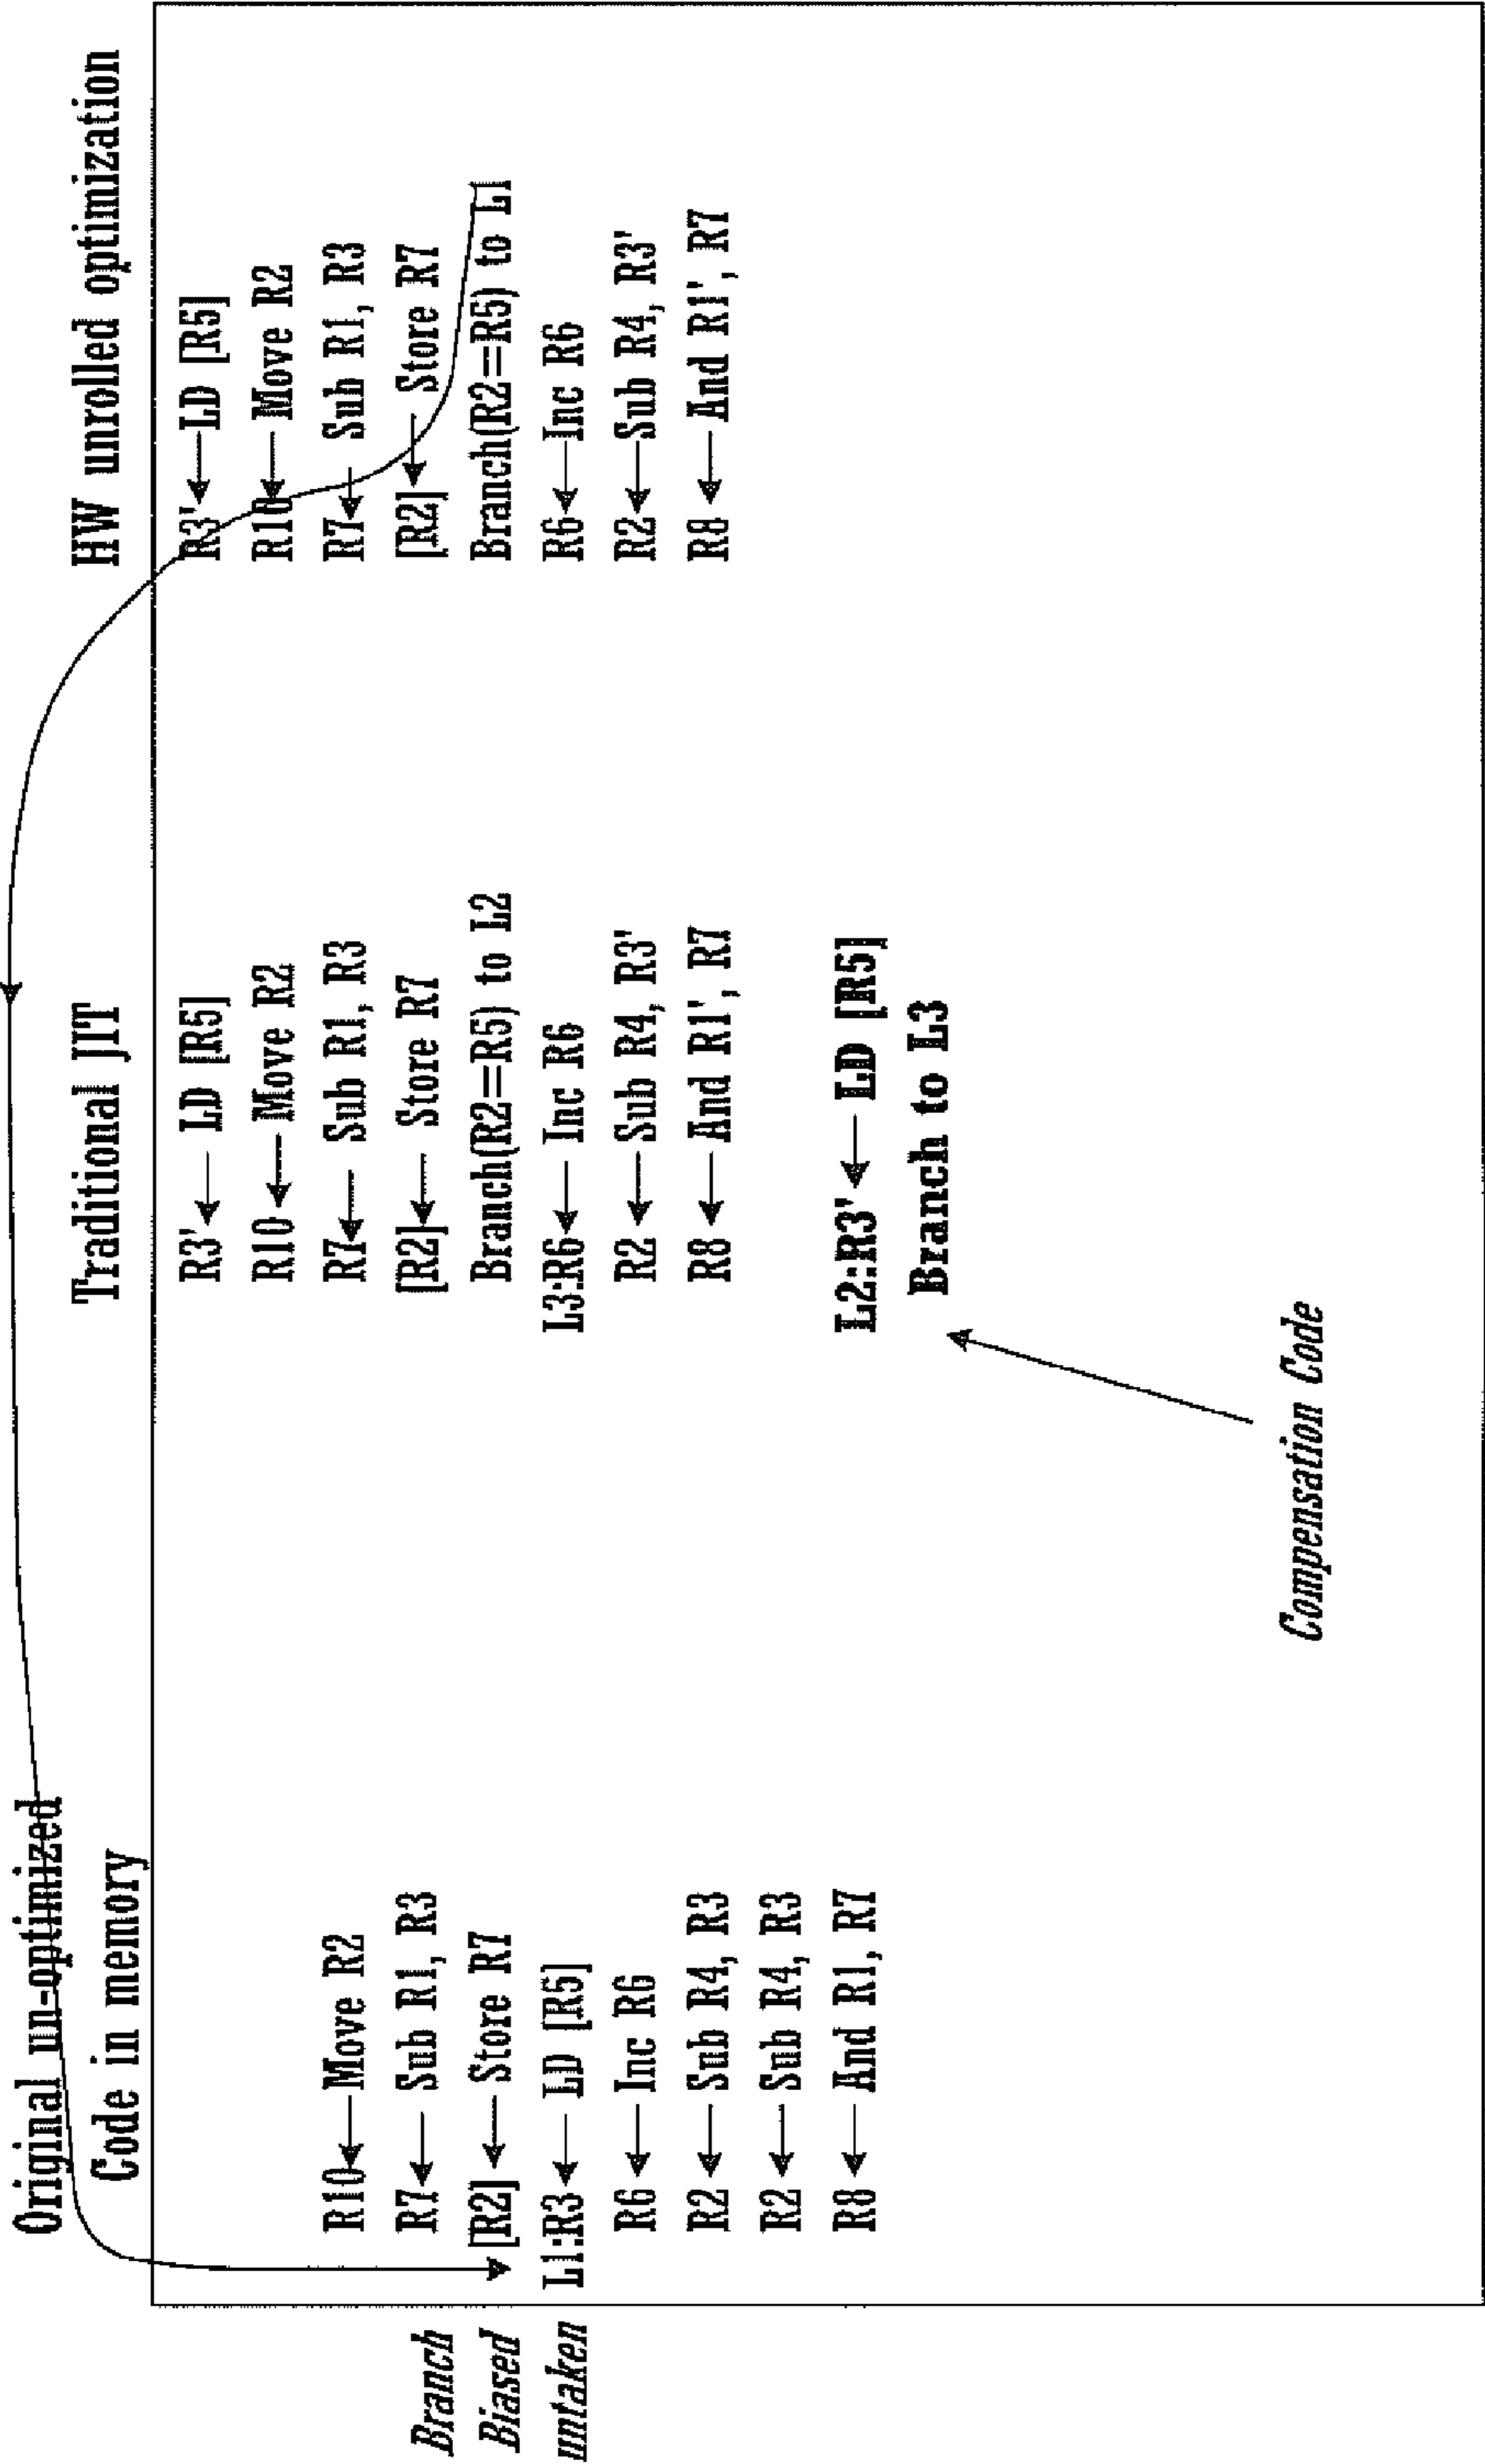


FIGURE 9

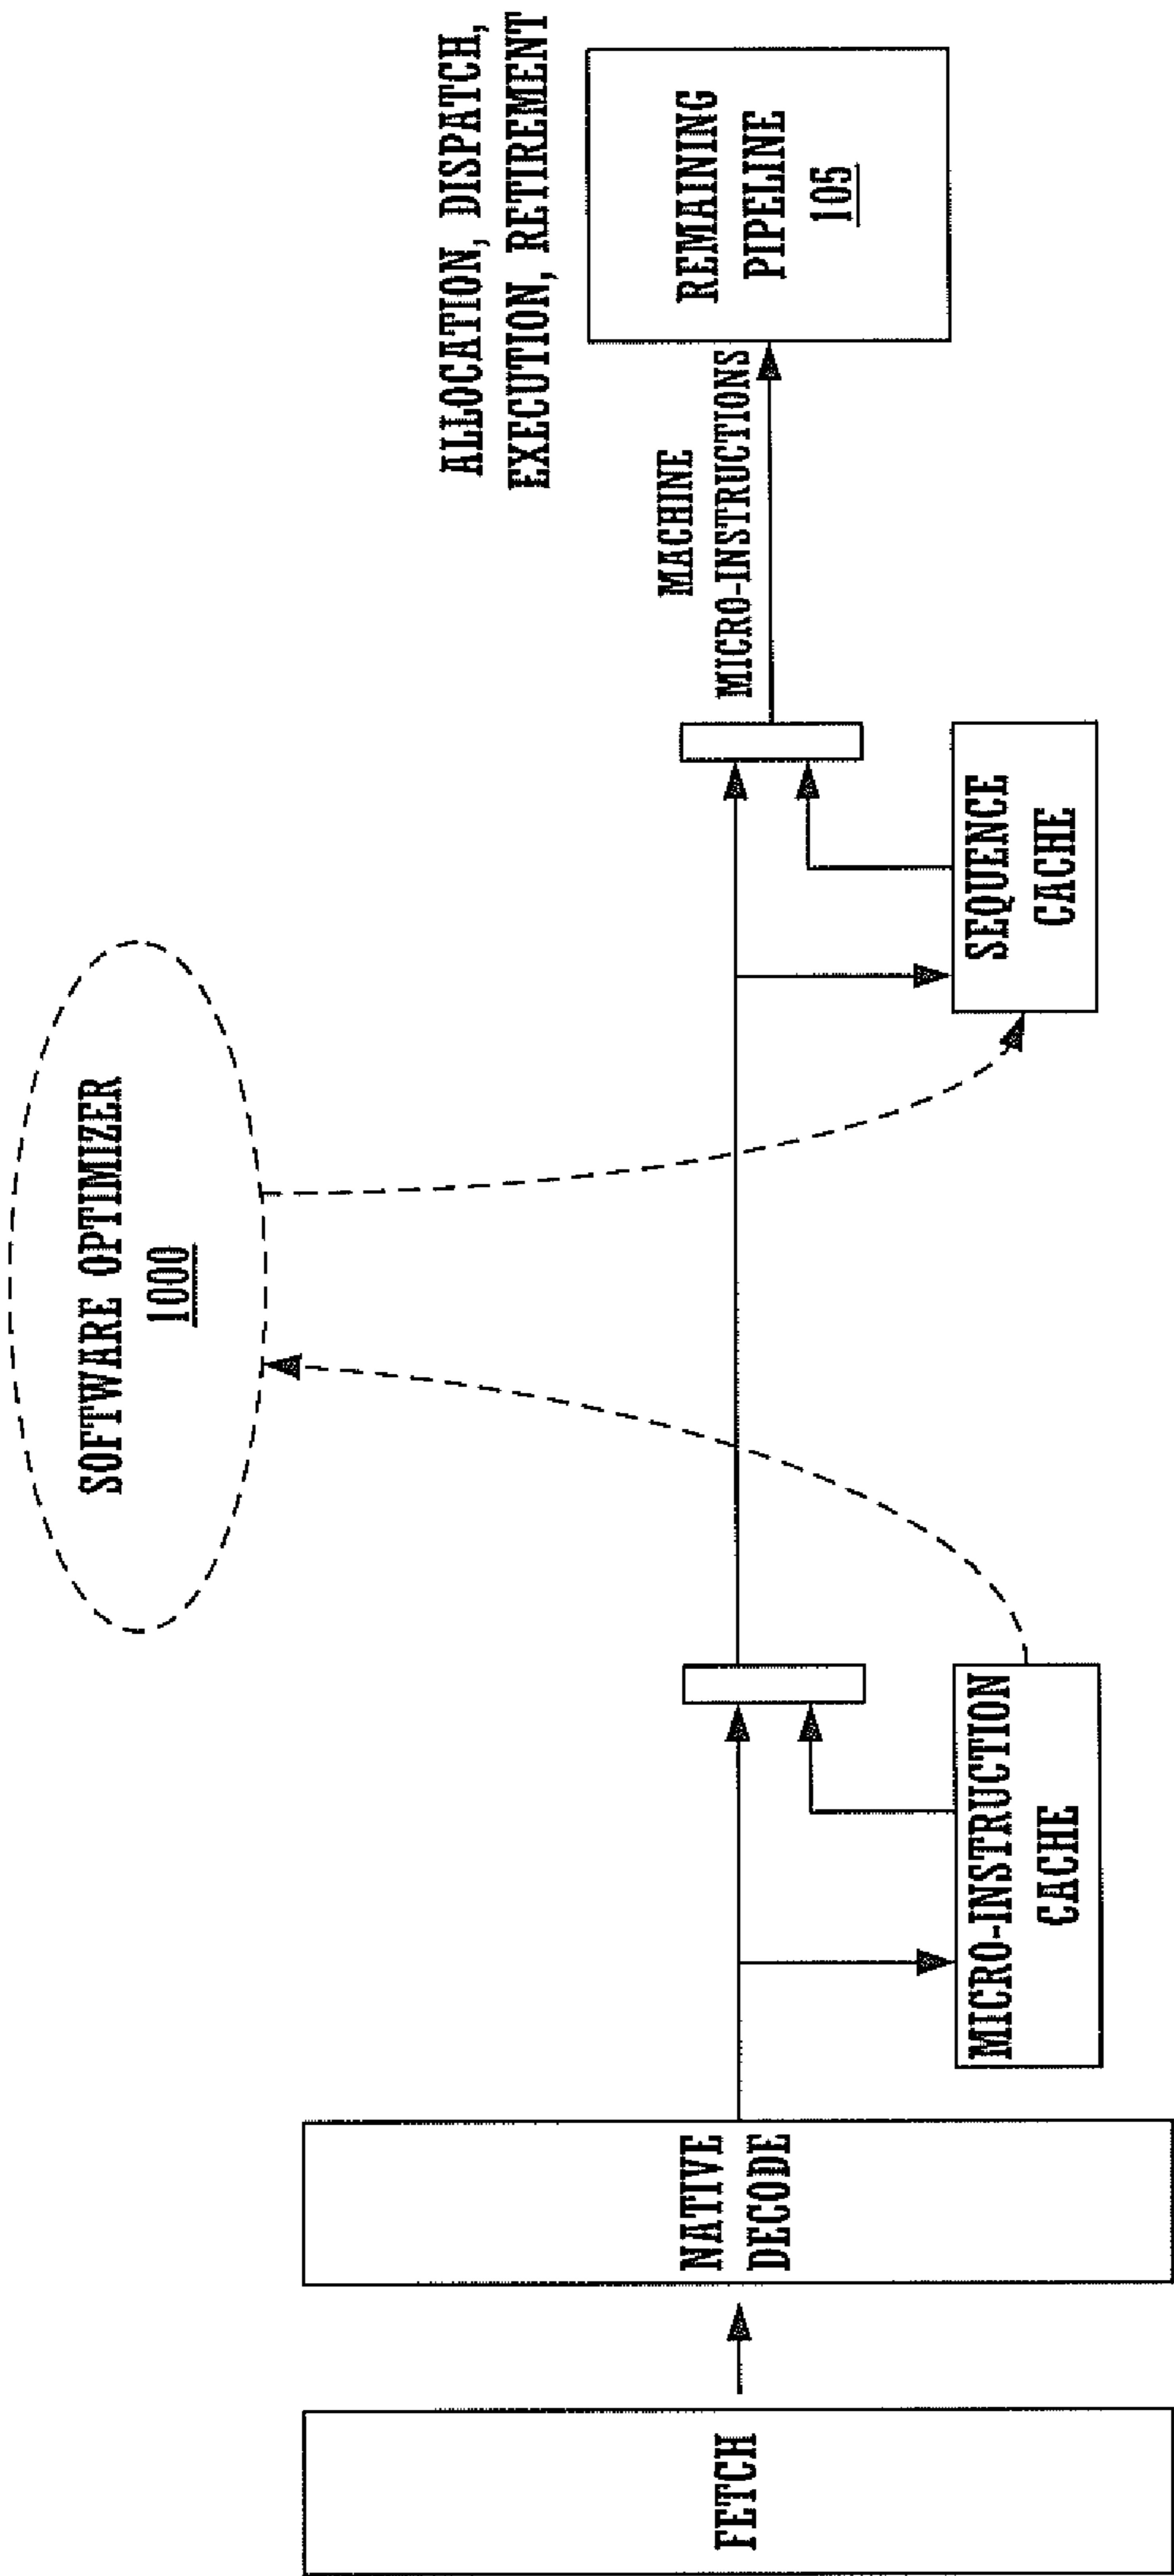


FIGURE 10

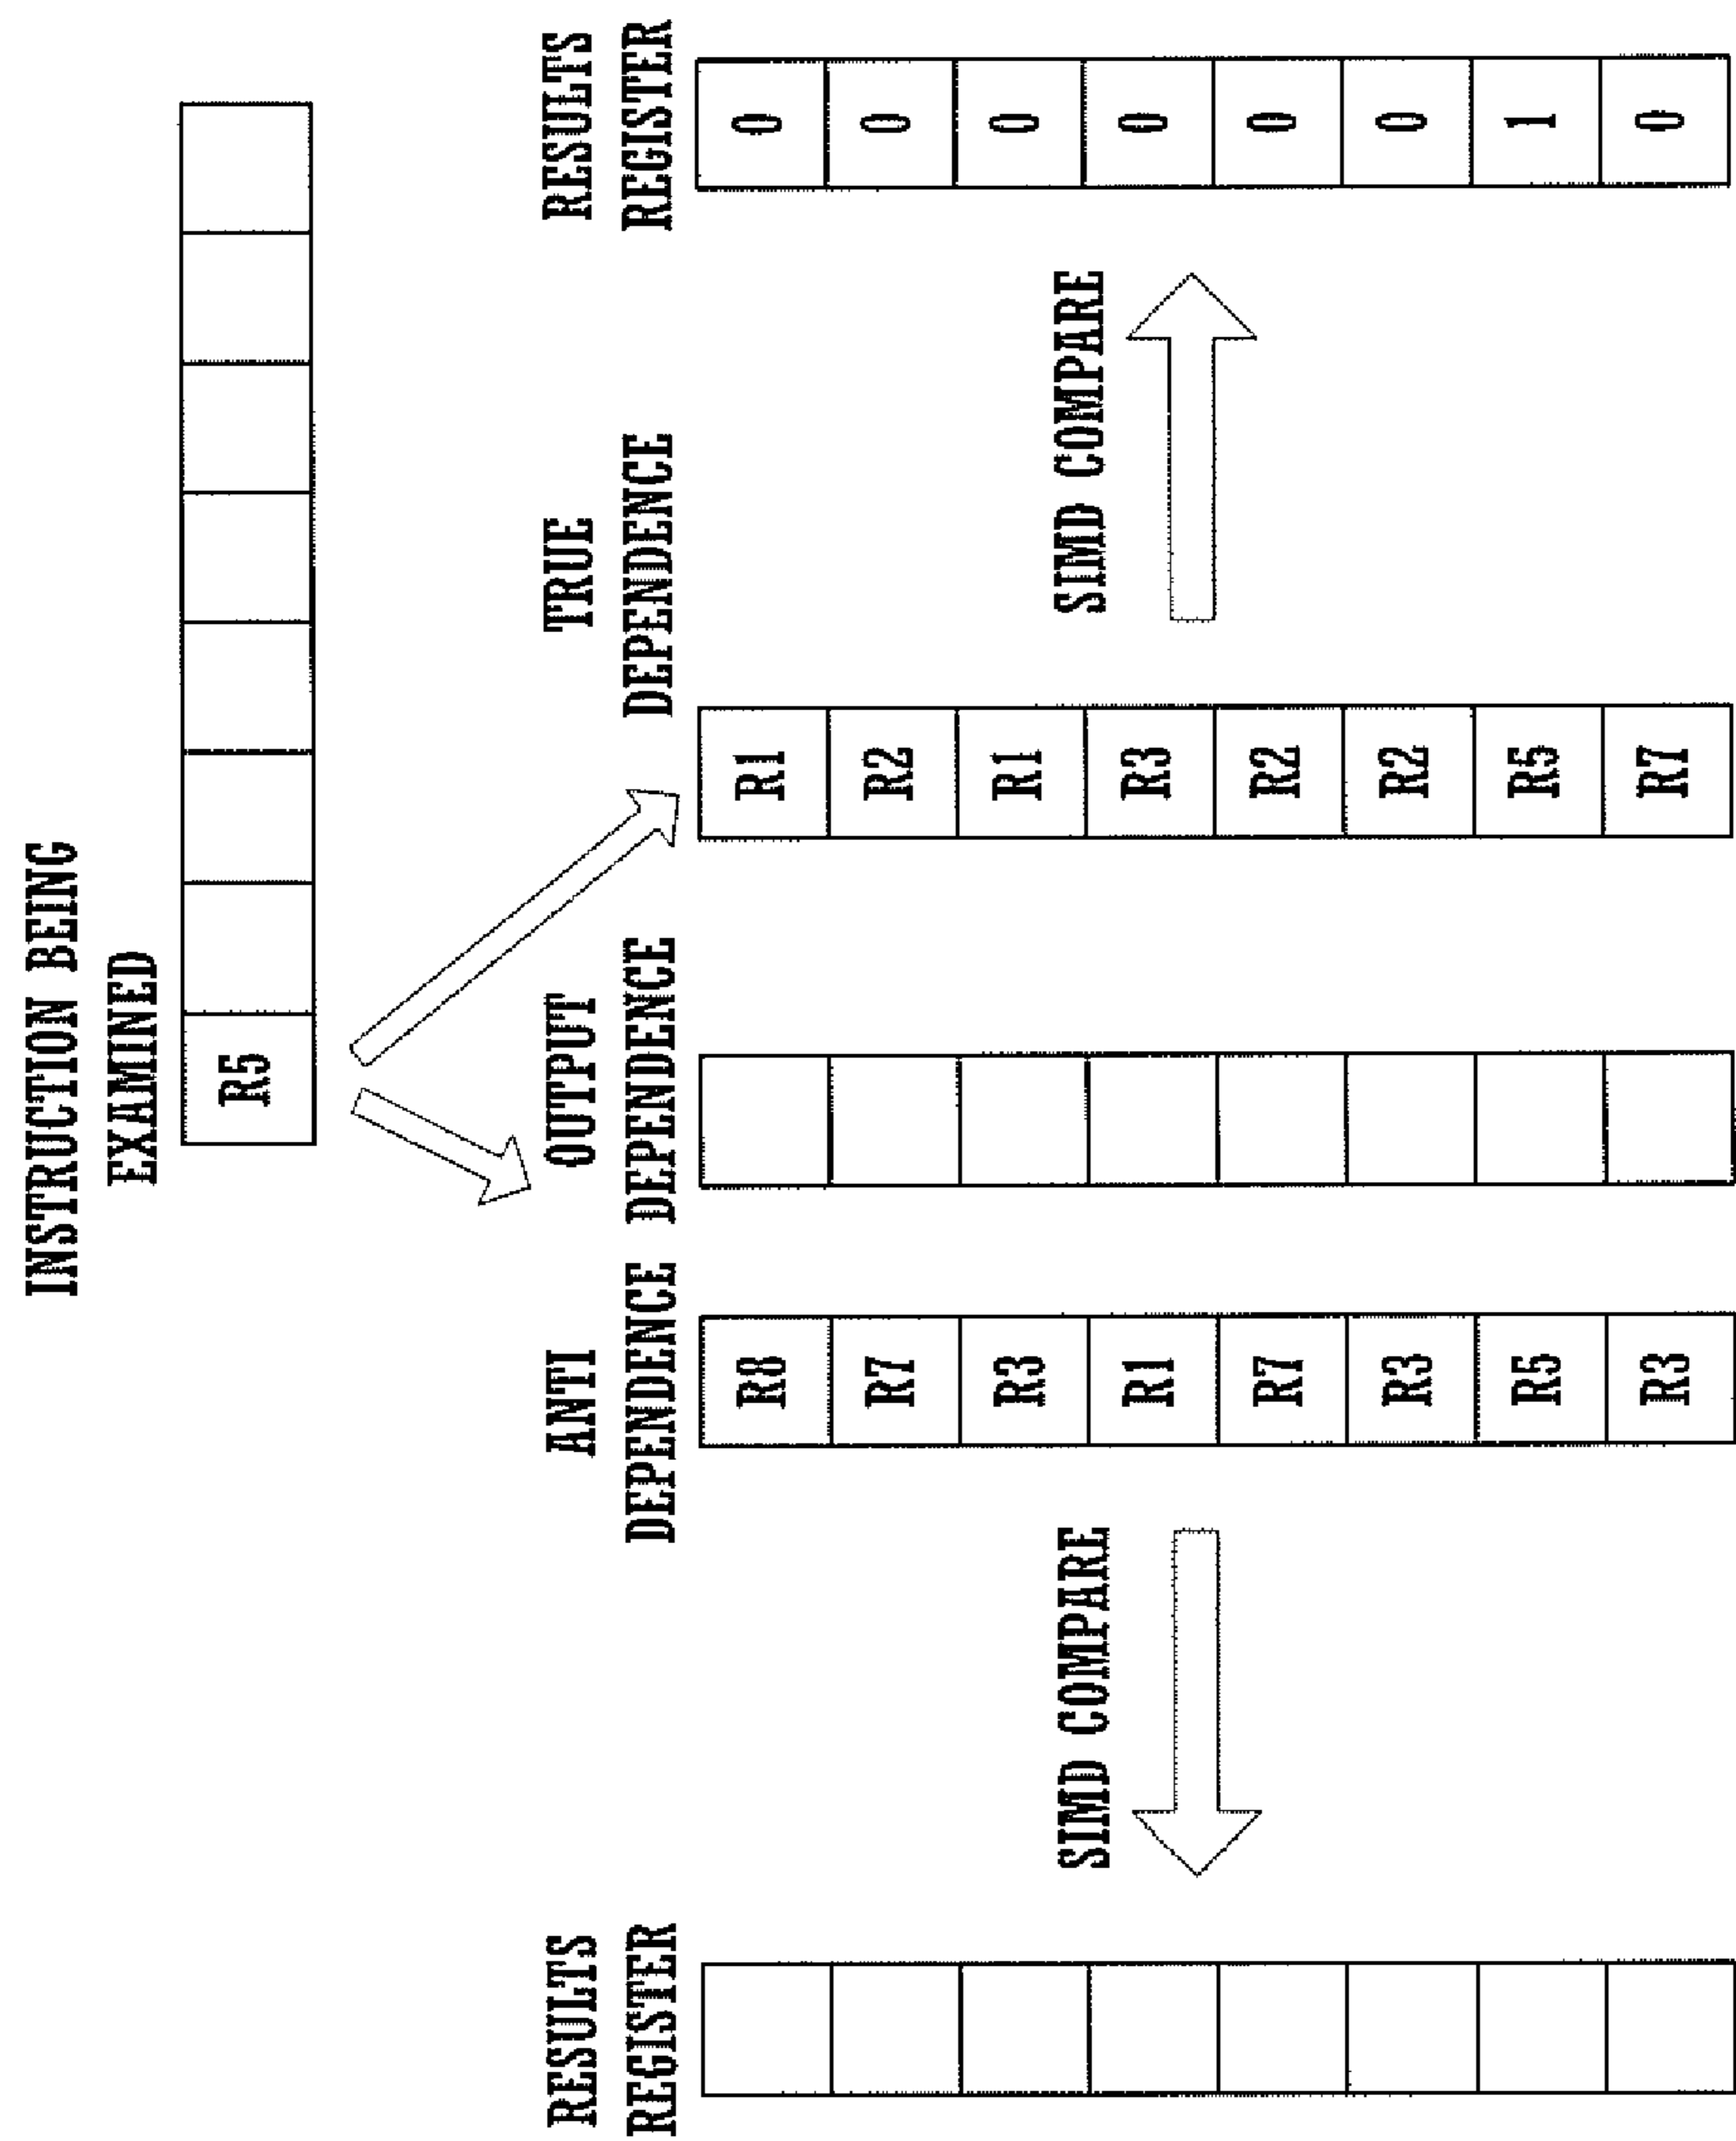
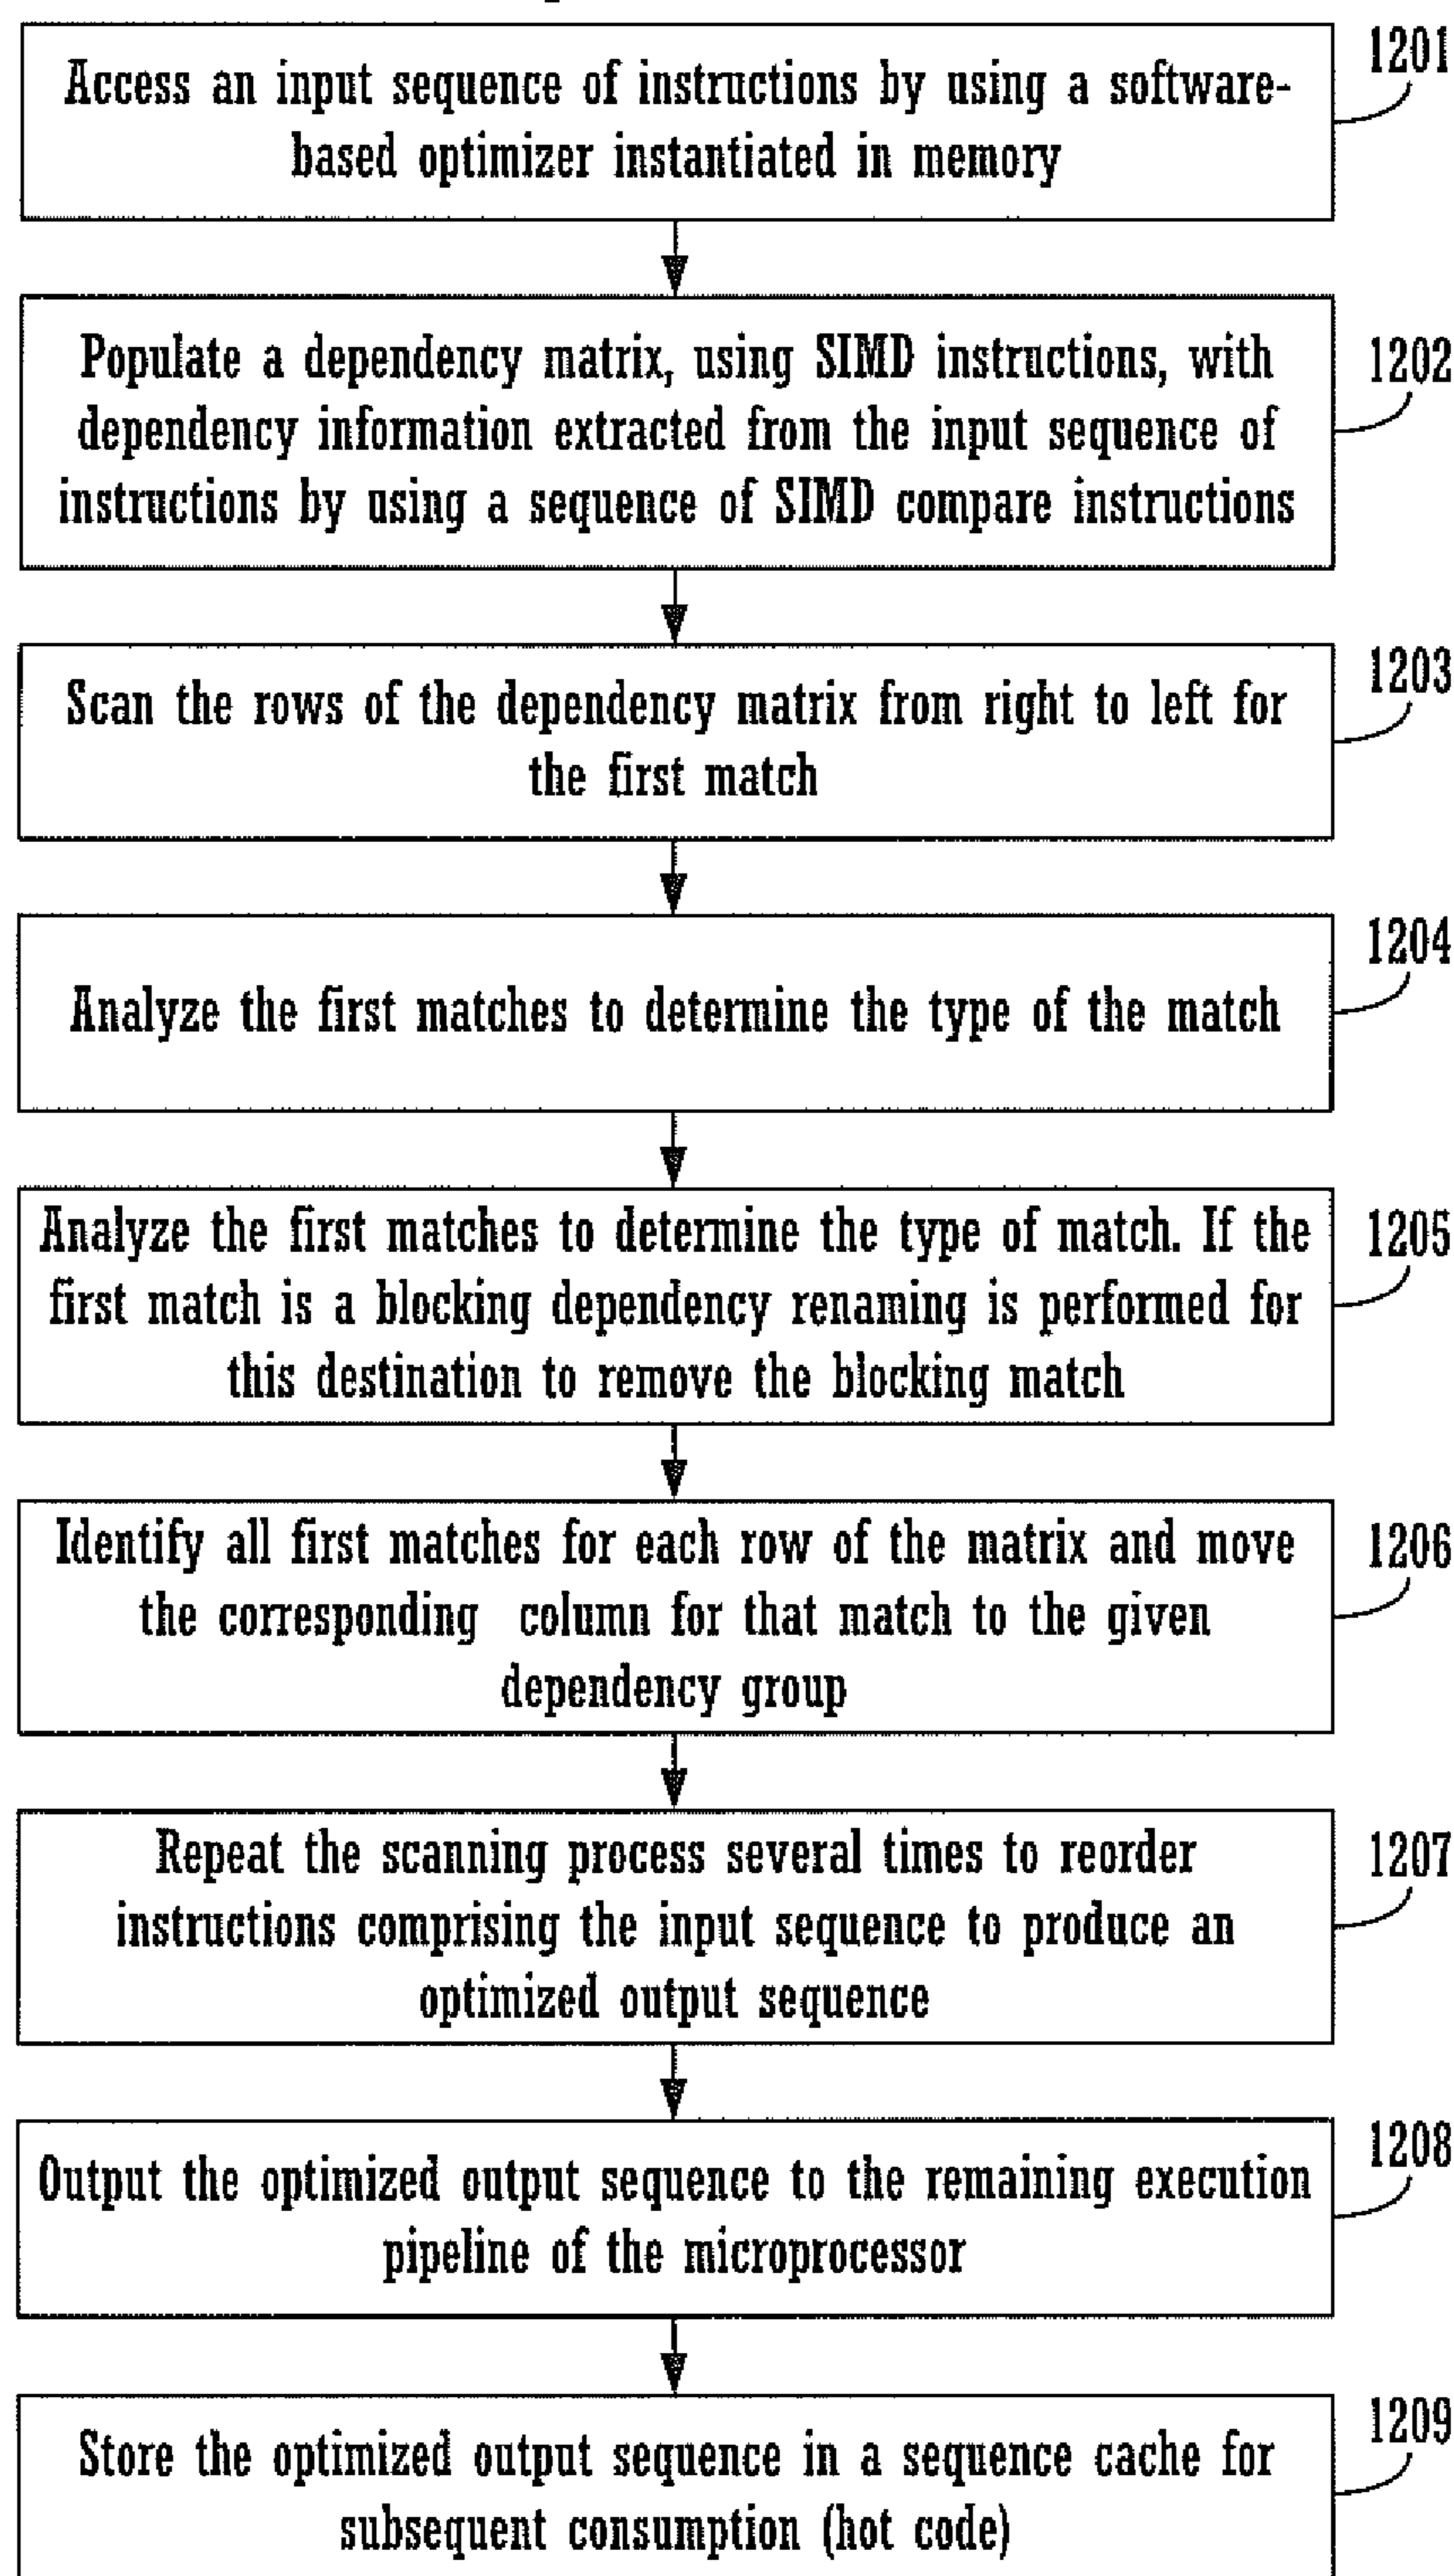


FIGURE 11

1200**Software Optimization Flowchart****FIGURE 12**

Dependency Broadcast (group & local)

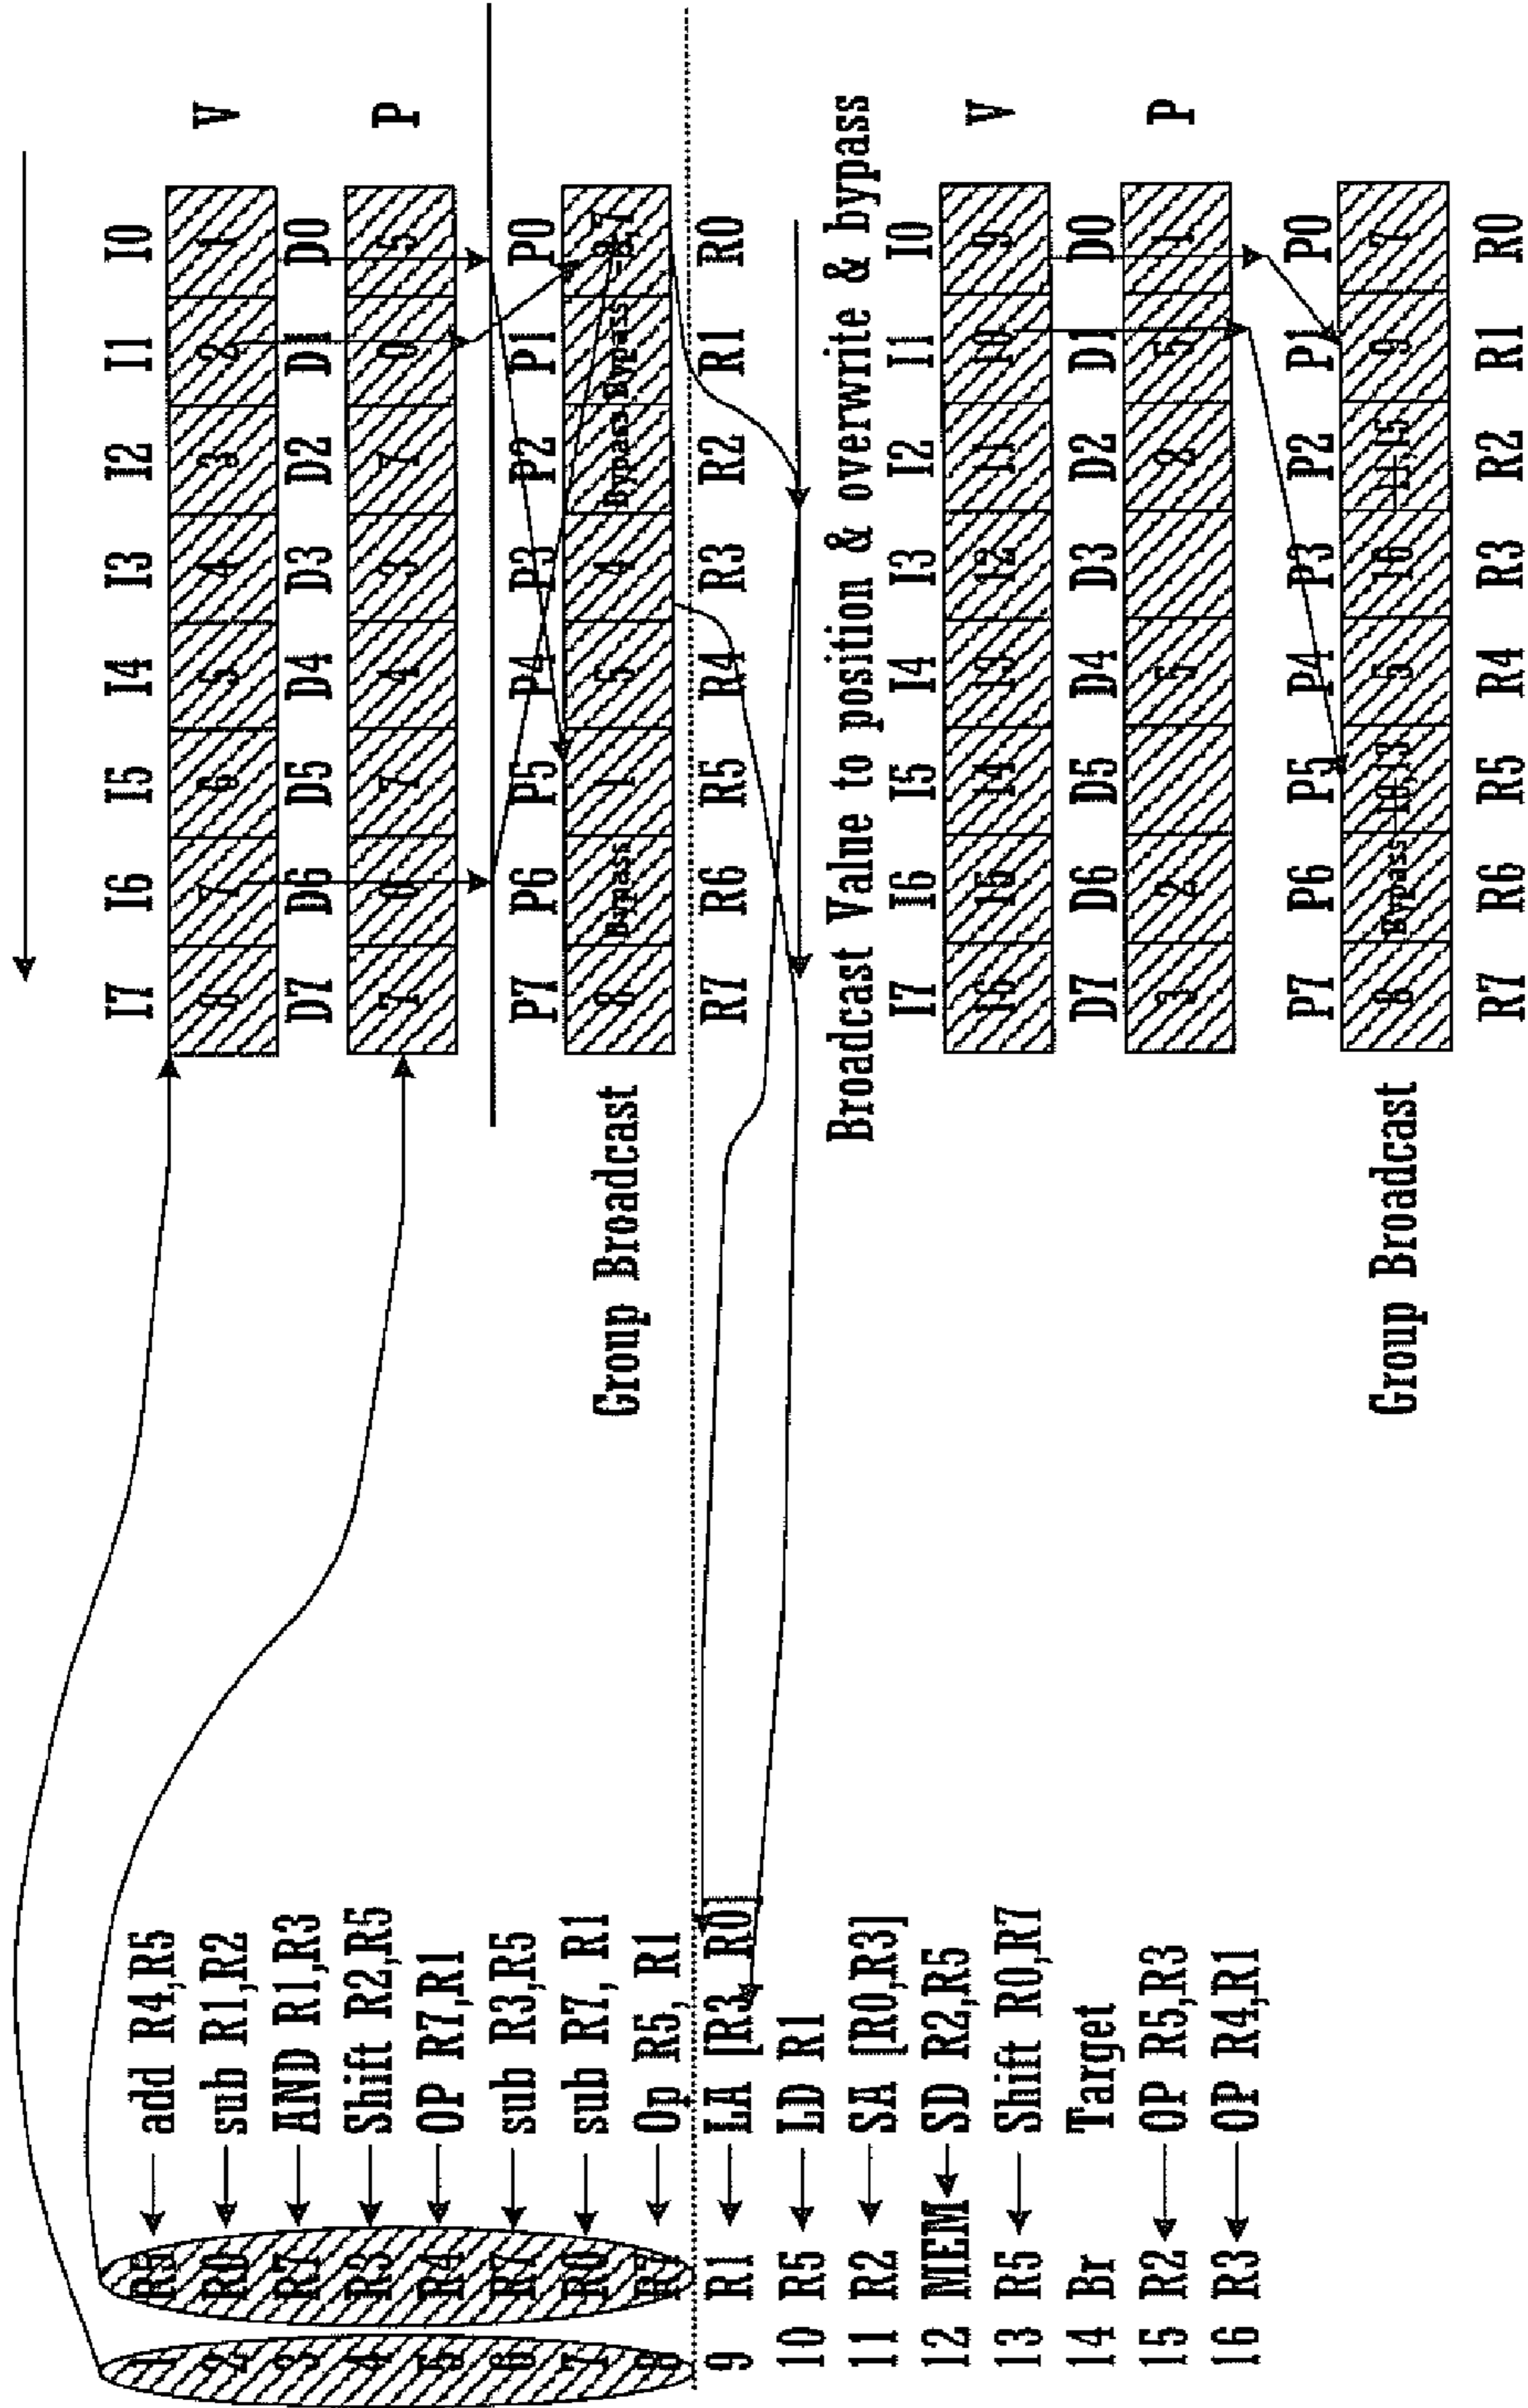


FIGURE 13

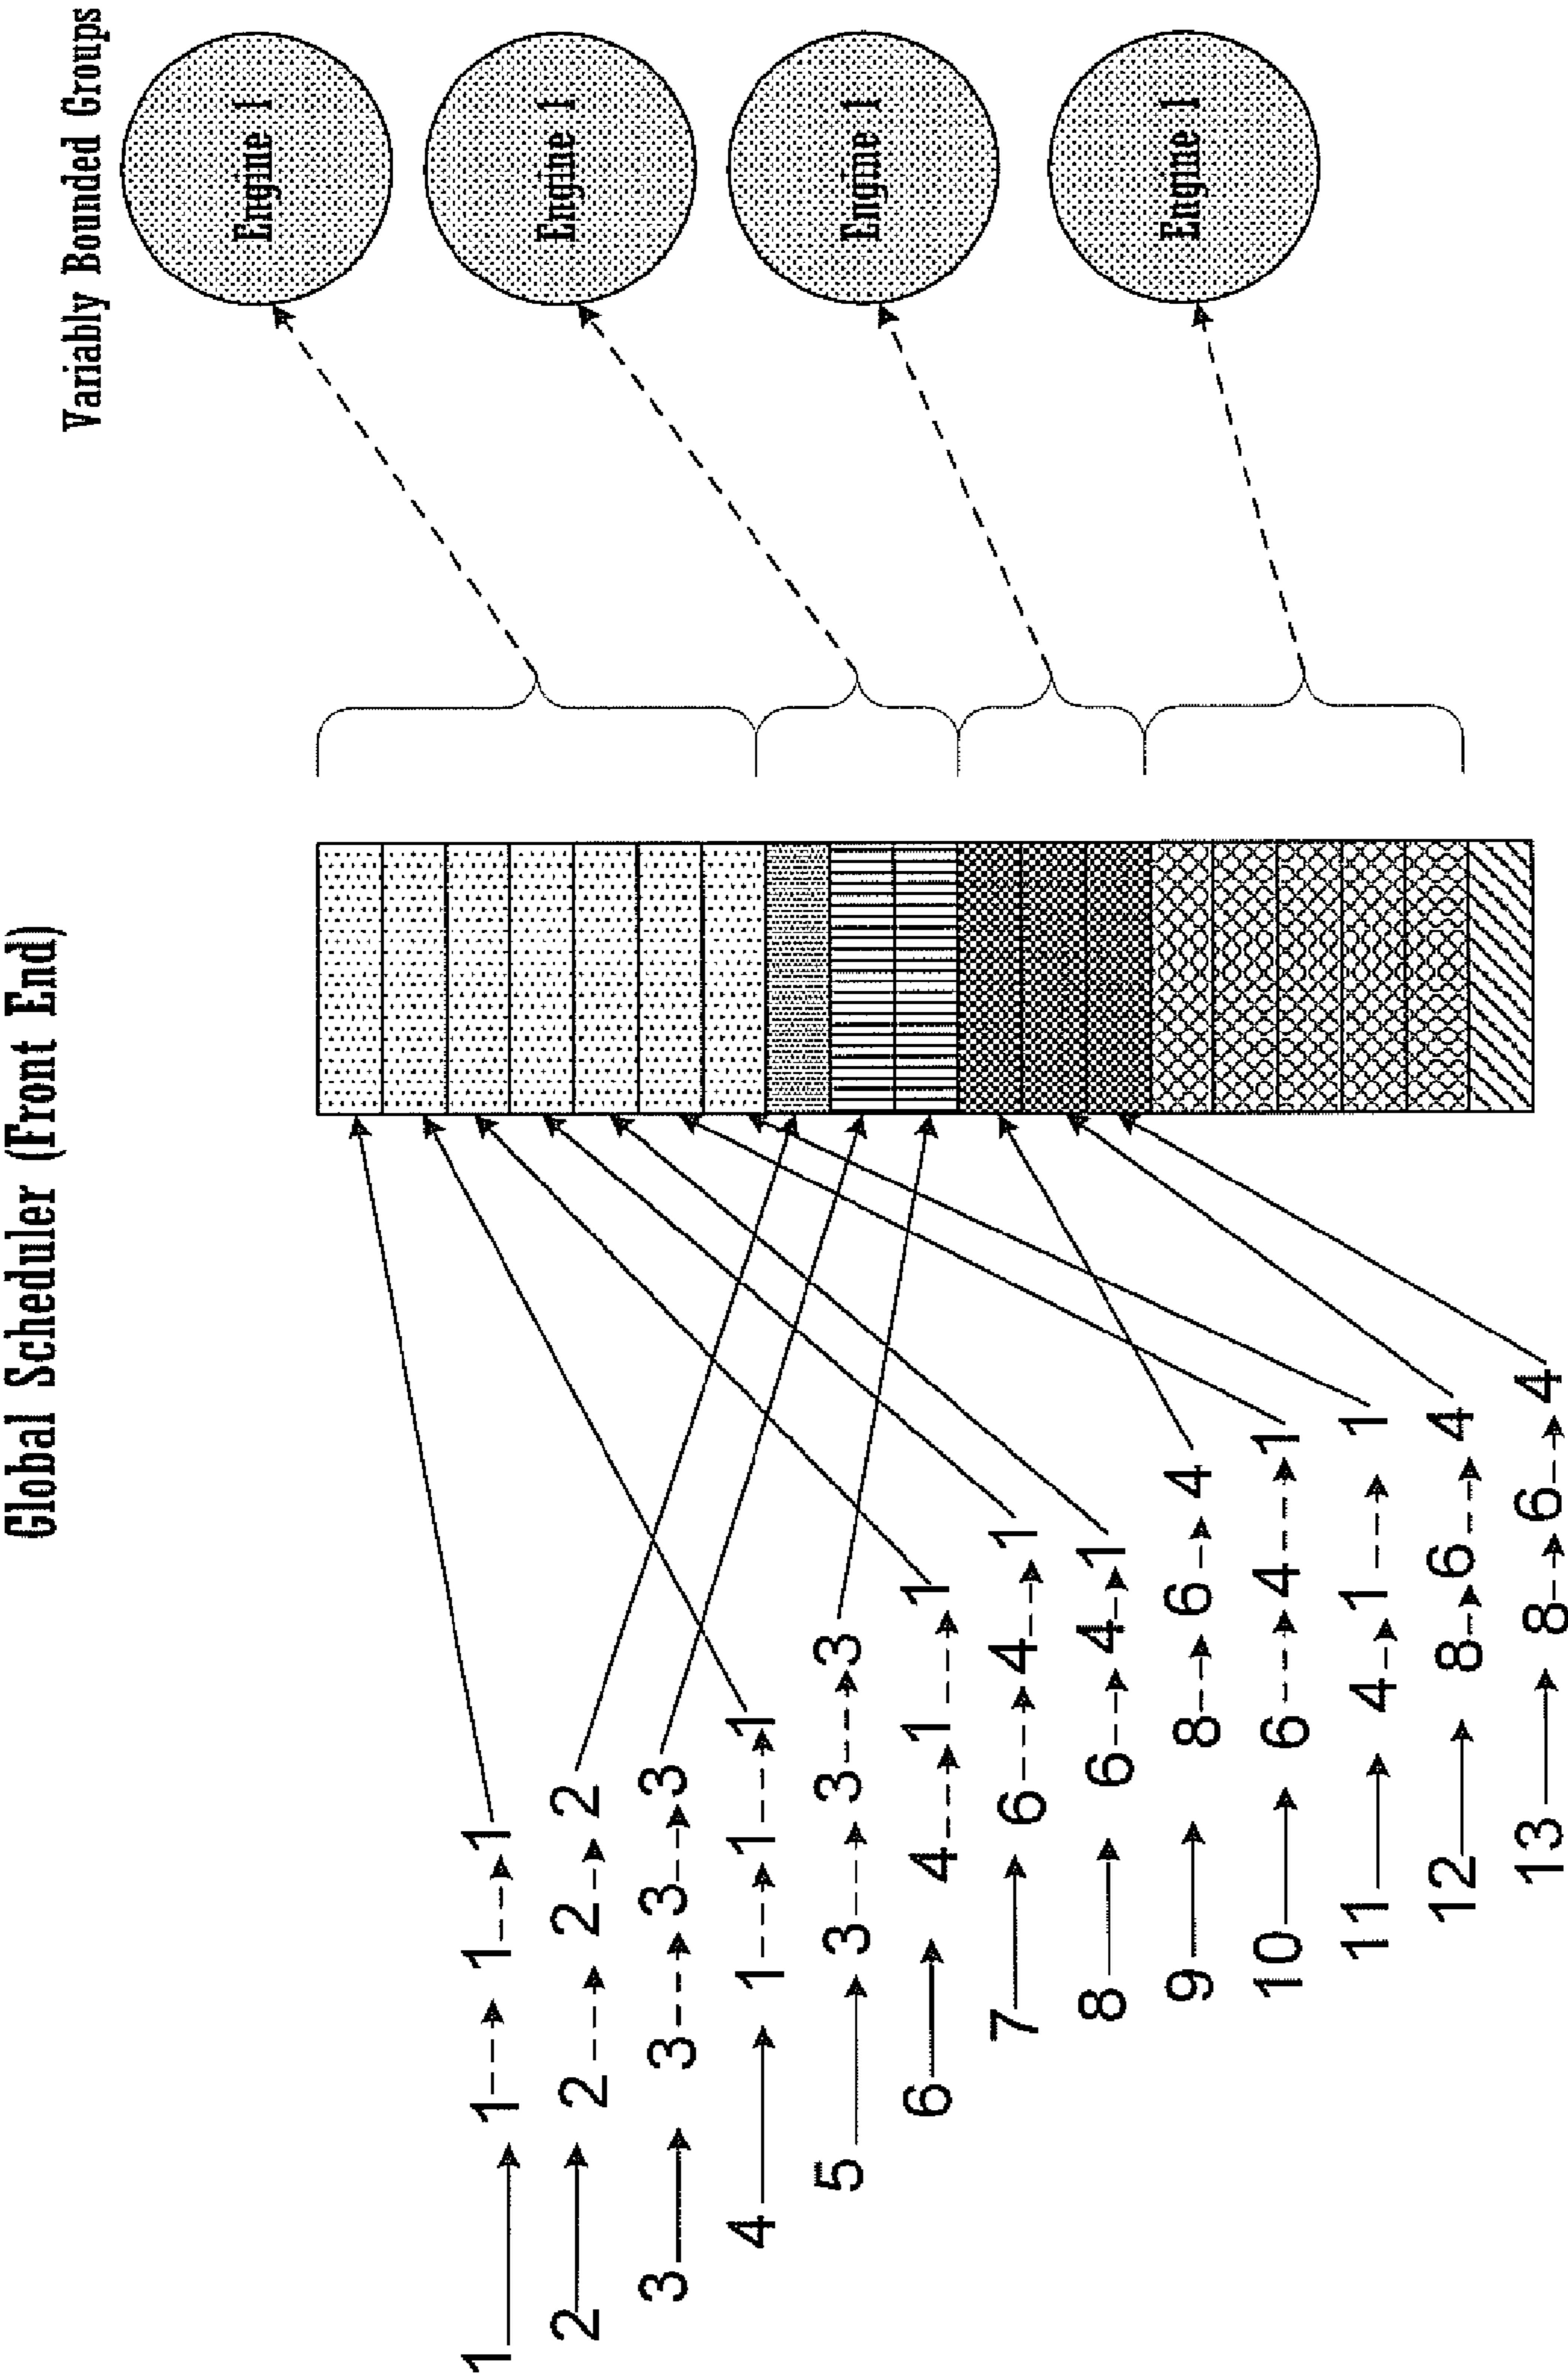


FIGURE 14

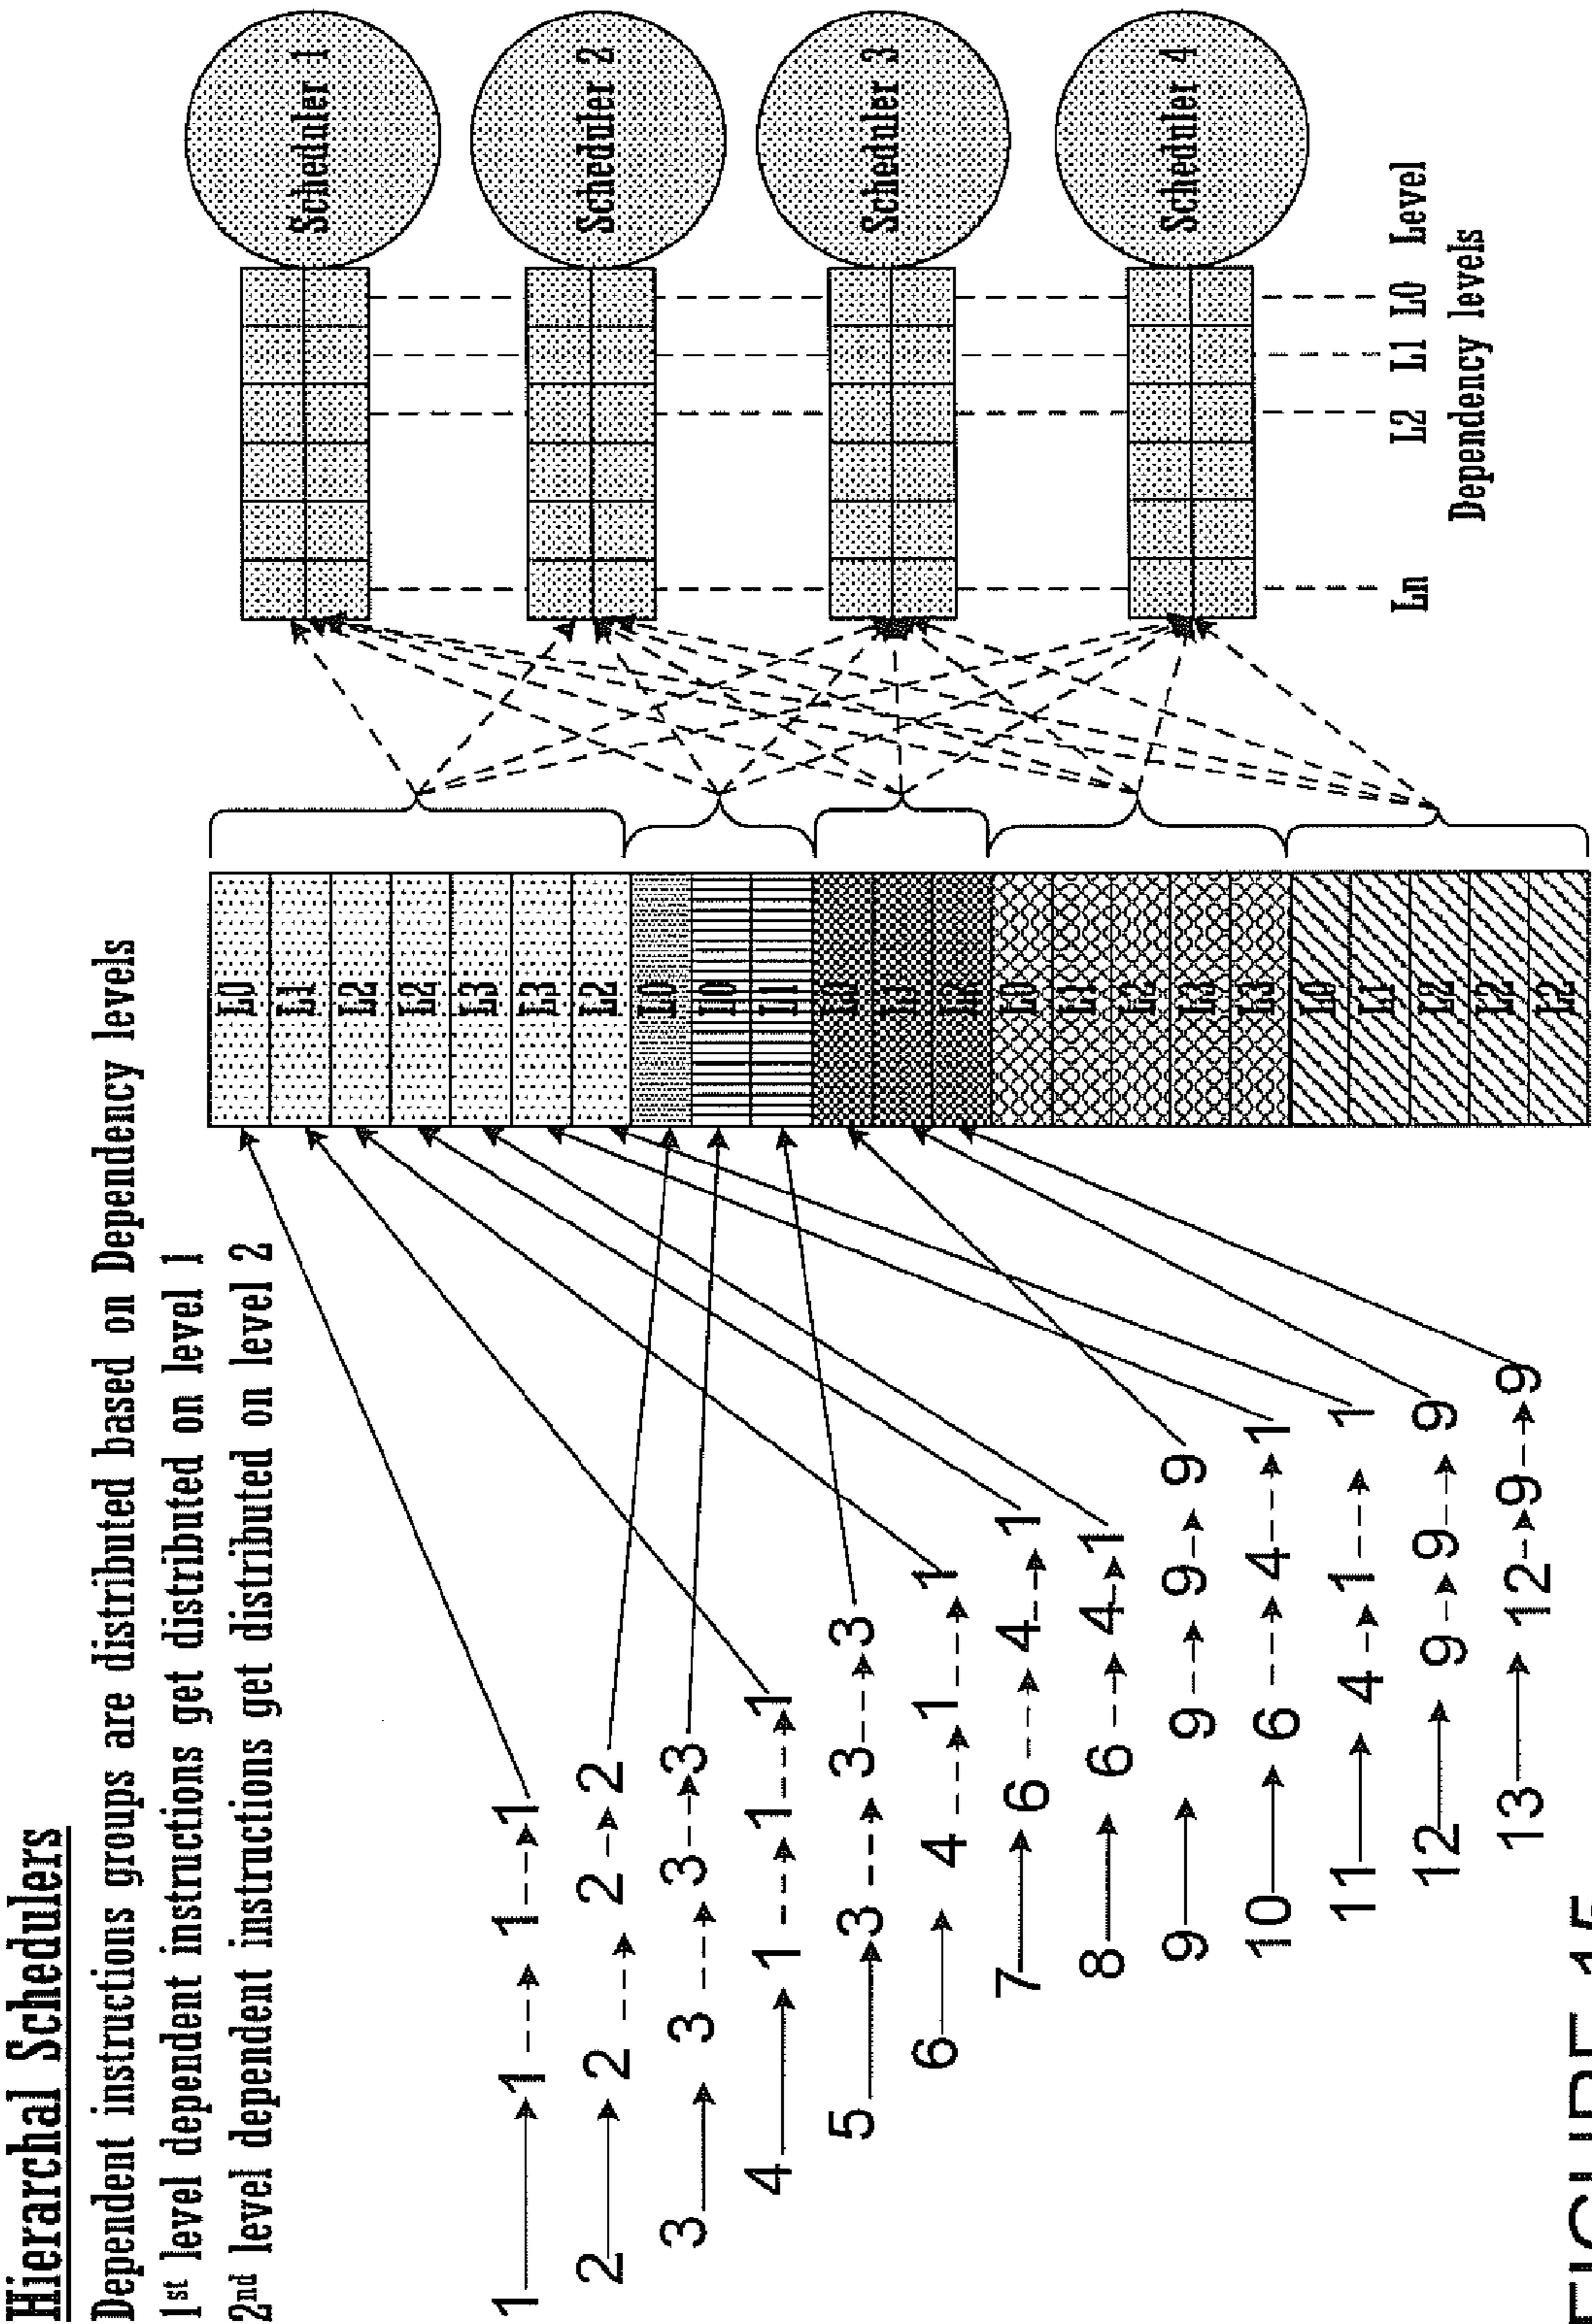
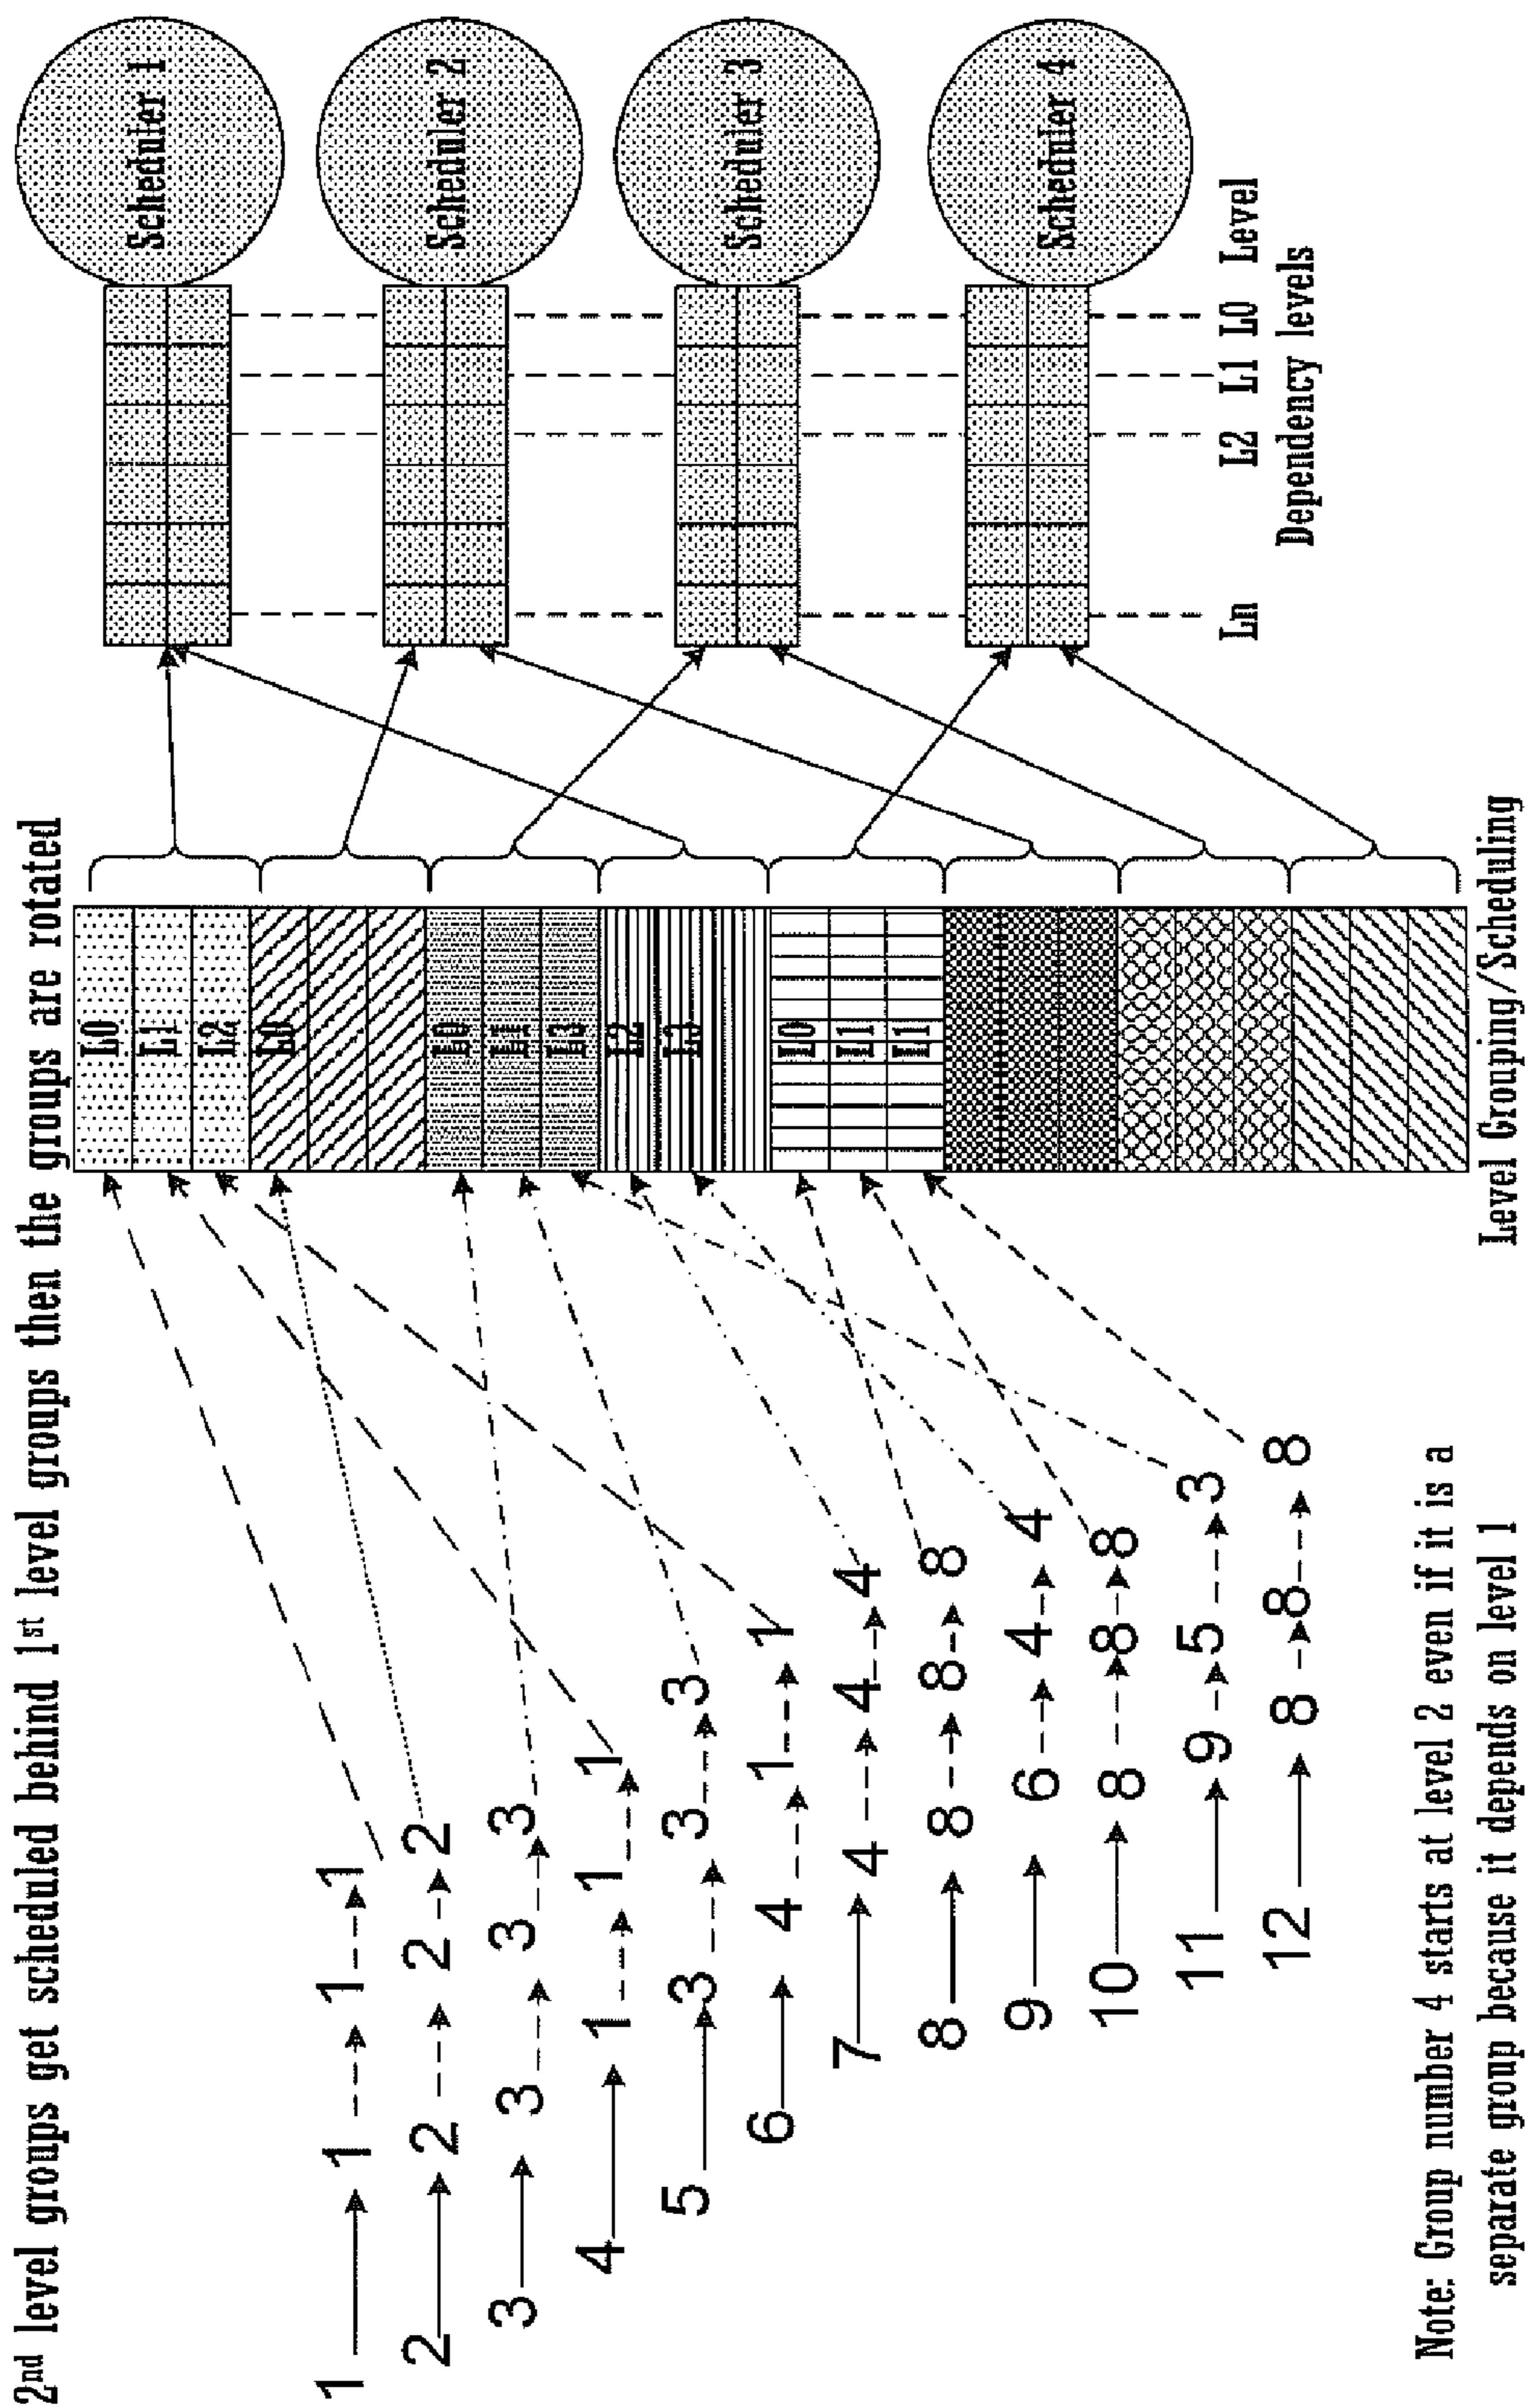


FIGURE 15

3 Slot Dependency Groups

Every 3 dependent instructions are scheduled on a scheduler

2nd level groups get scheduled behind 1st level groups then the groups are rotated



Note: Group number 4 starts at level 2 even if it is a separate group because it depends on level 1

Level Grouping/Scheduling

FIGURE 16

3 Slot Dependency Groups

Every 3 dependent instructions are scheduled on a scheduler

2nd level groups get scheduled behind 1st level groups then the groups are rotated

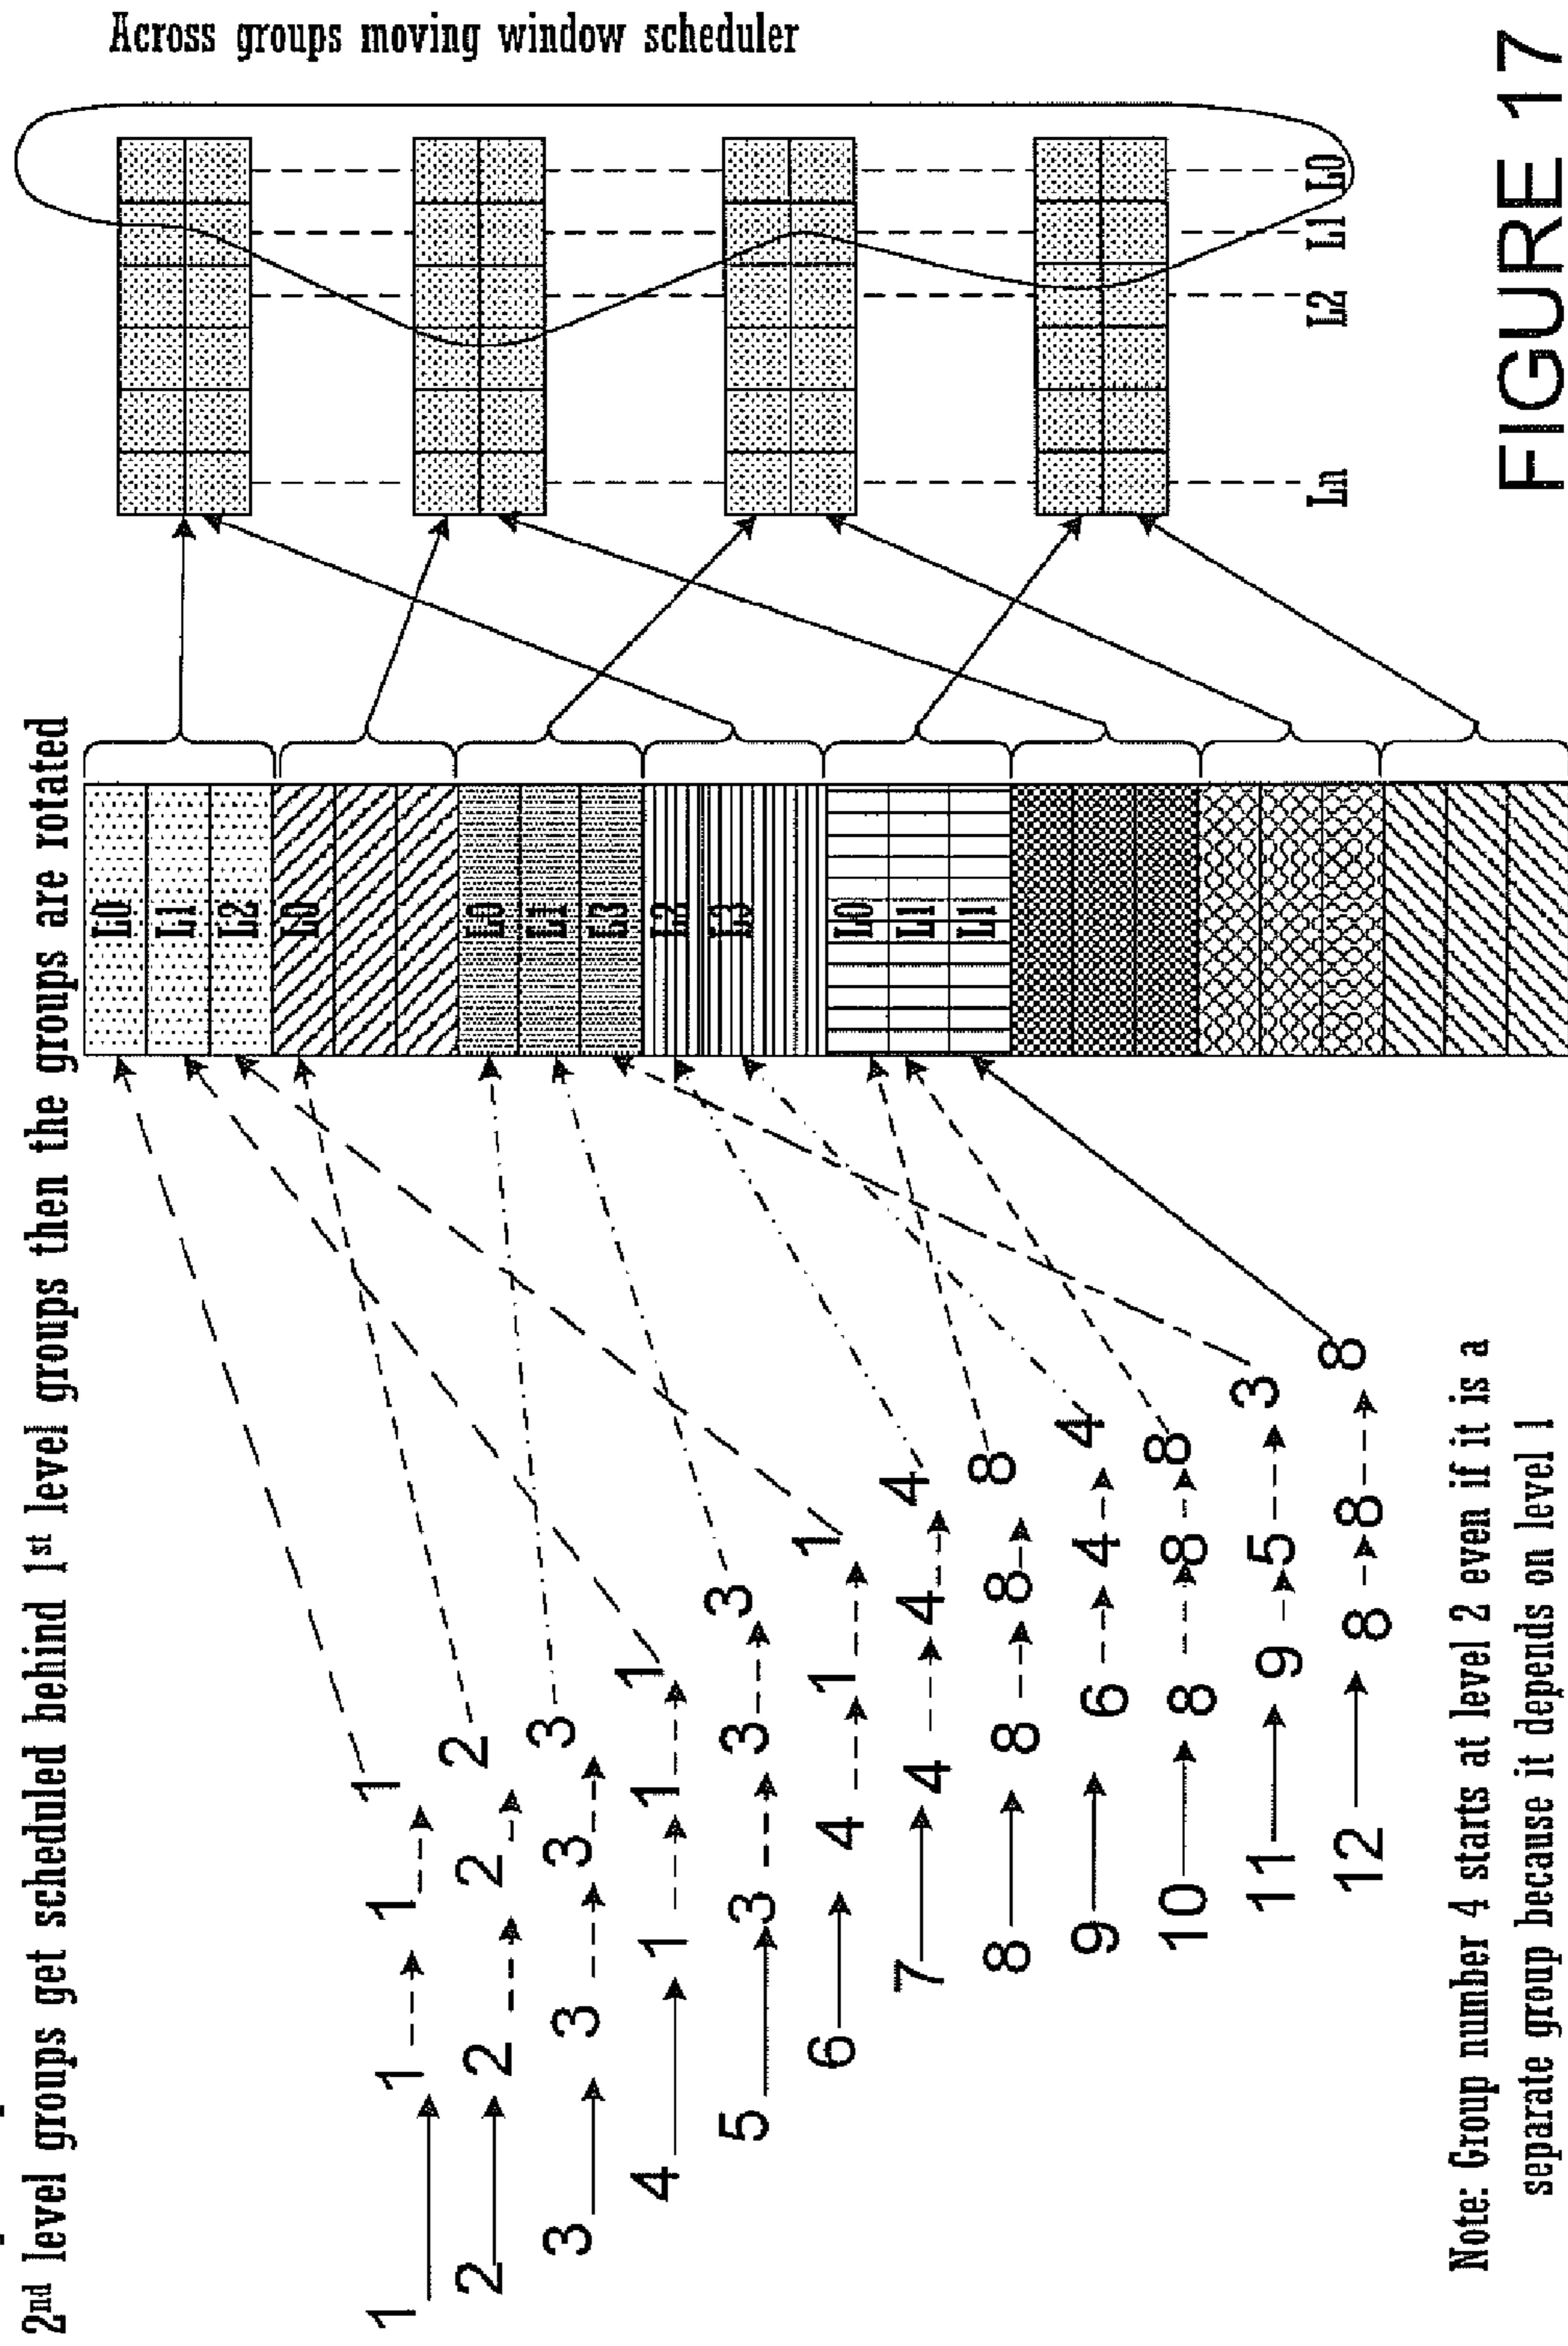


FIGURE 17

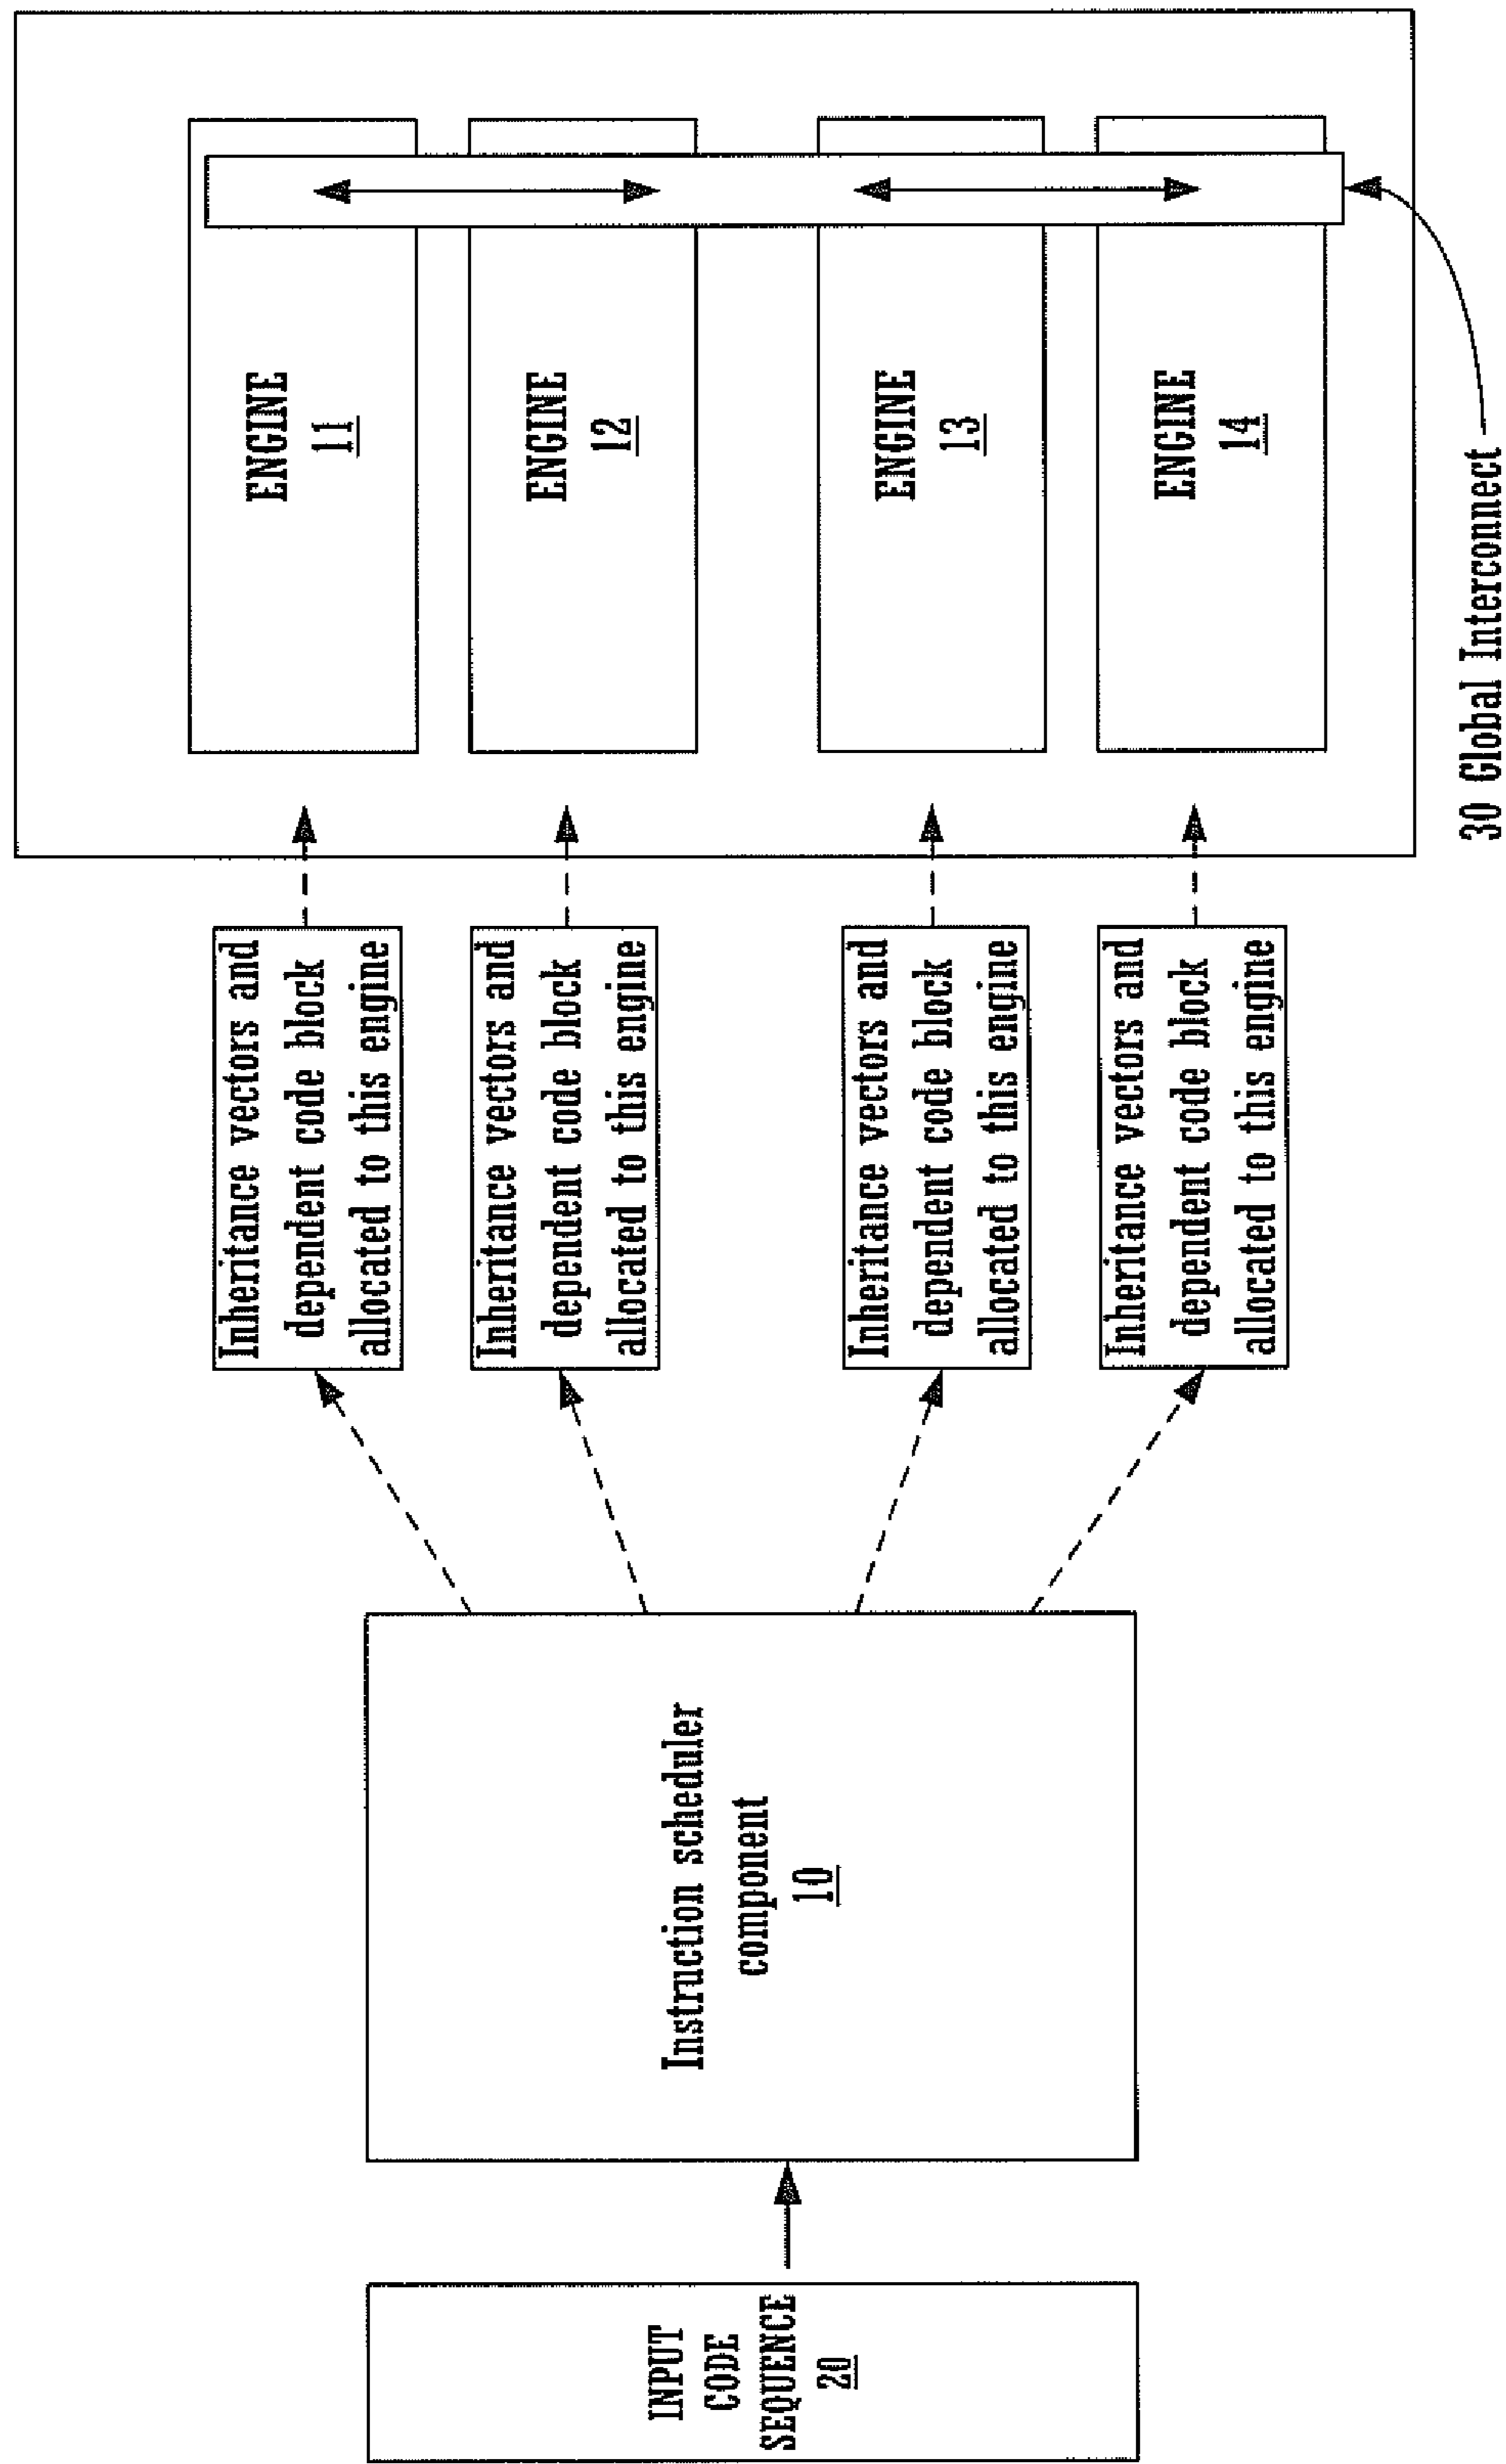


FIGURE 18

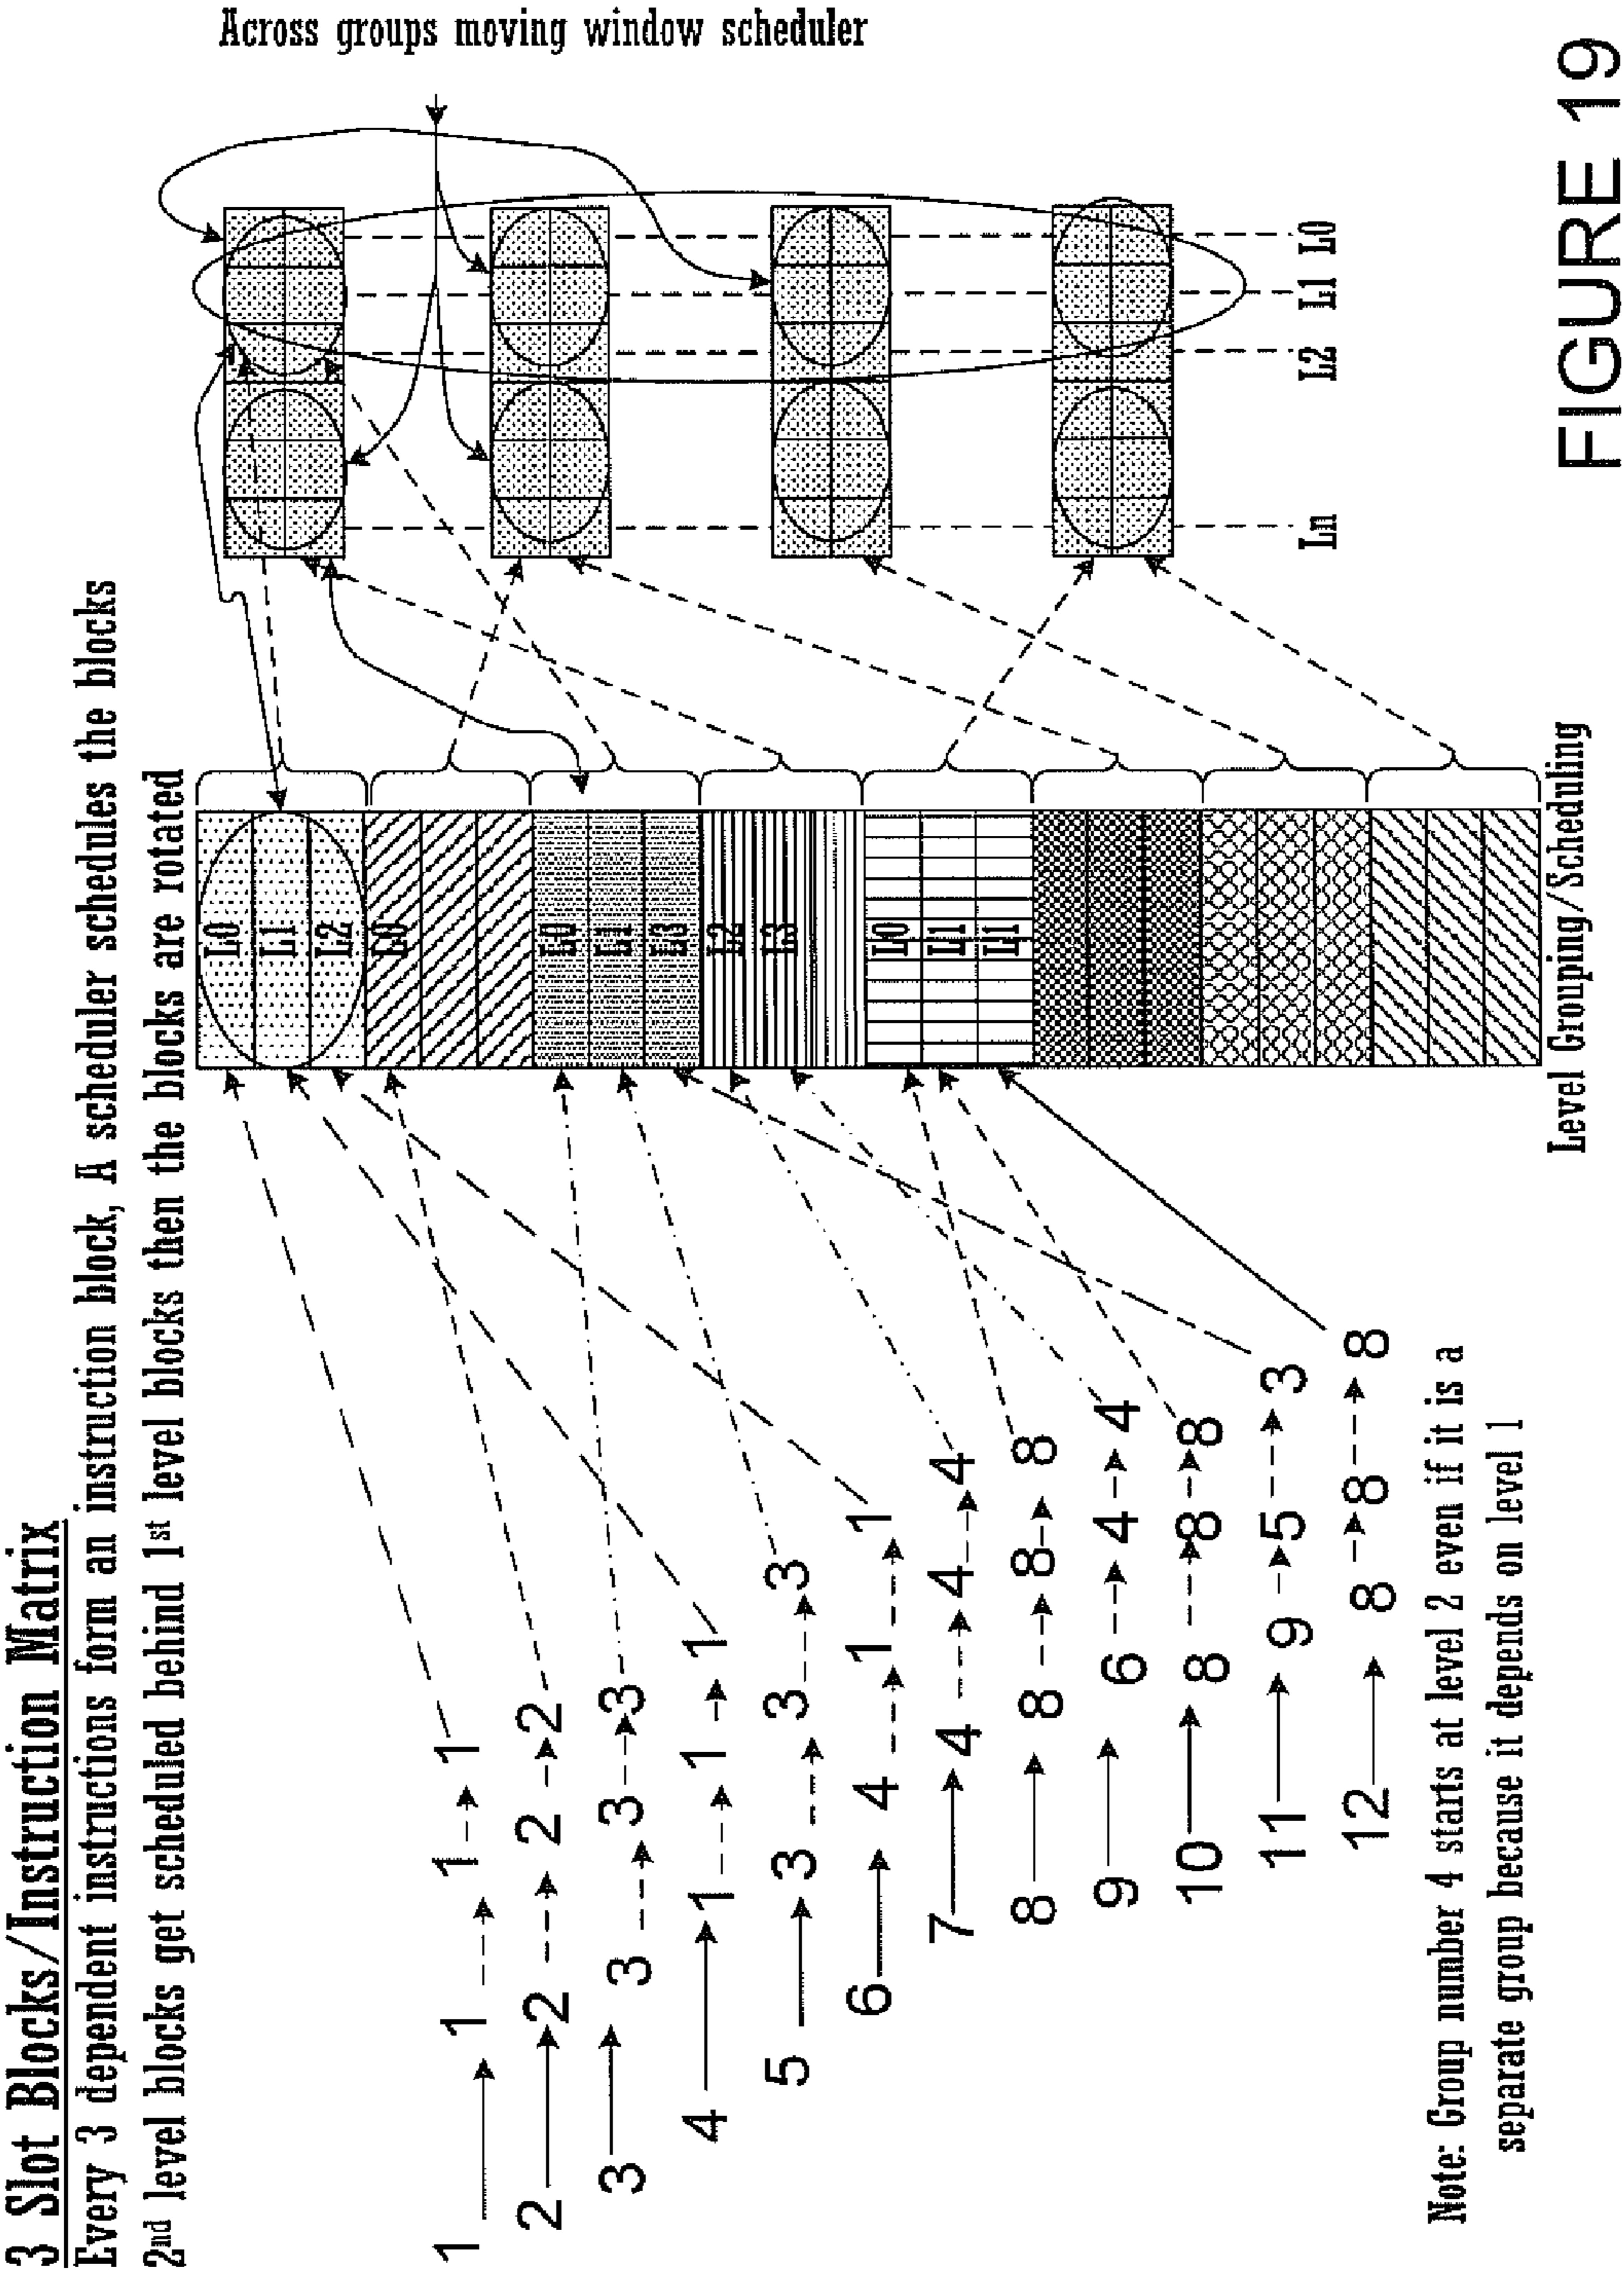


FIGURE 19

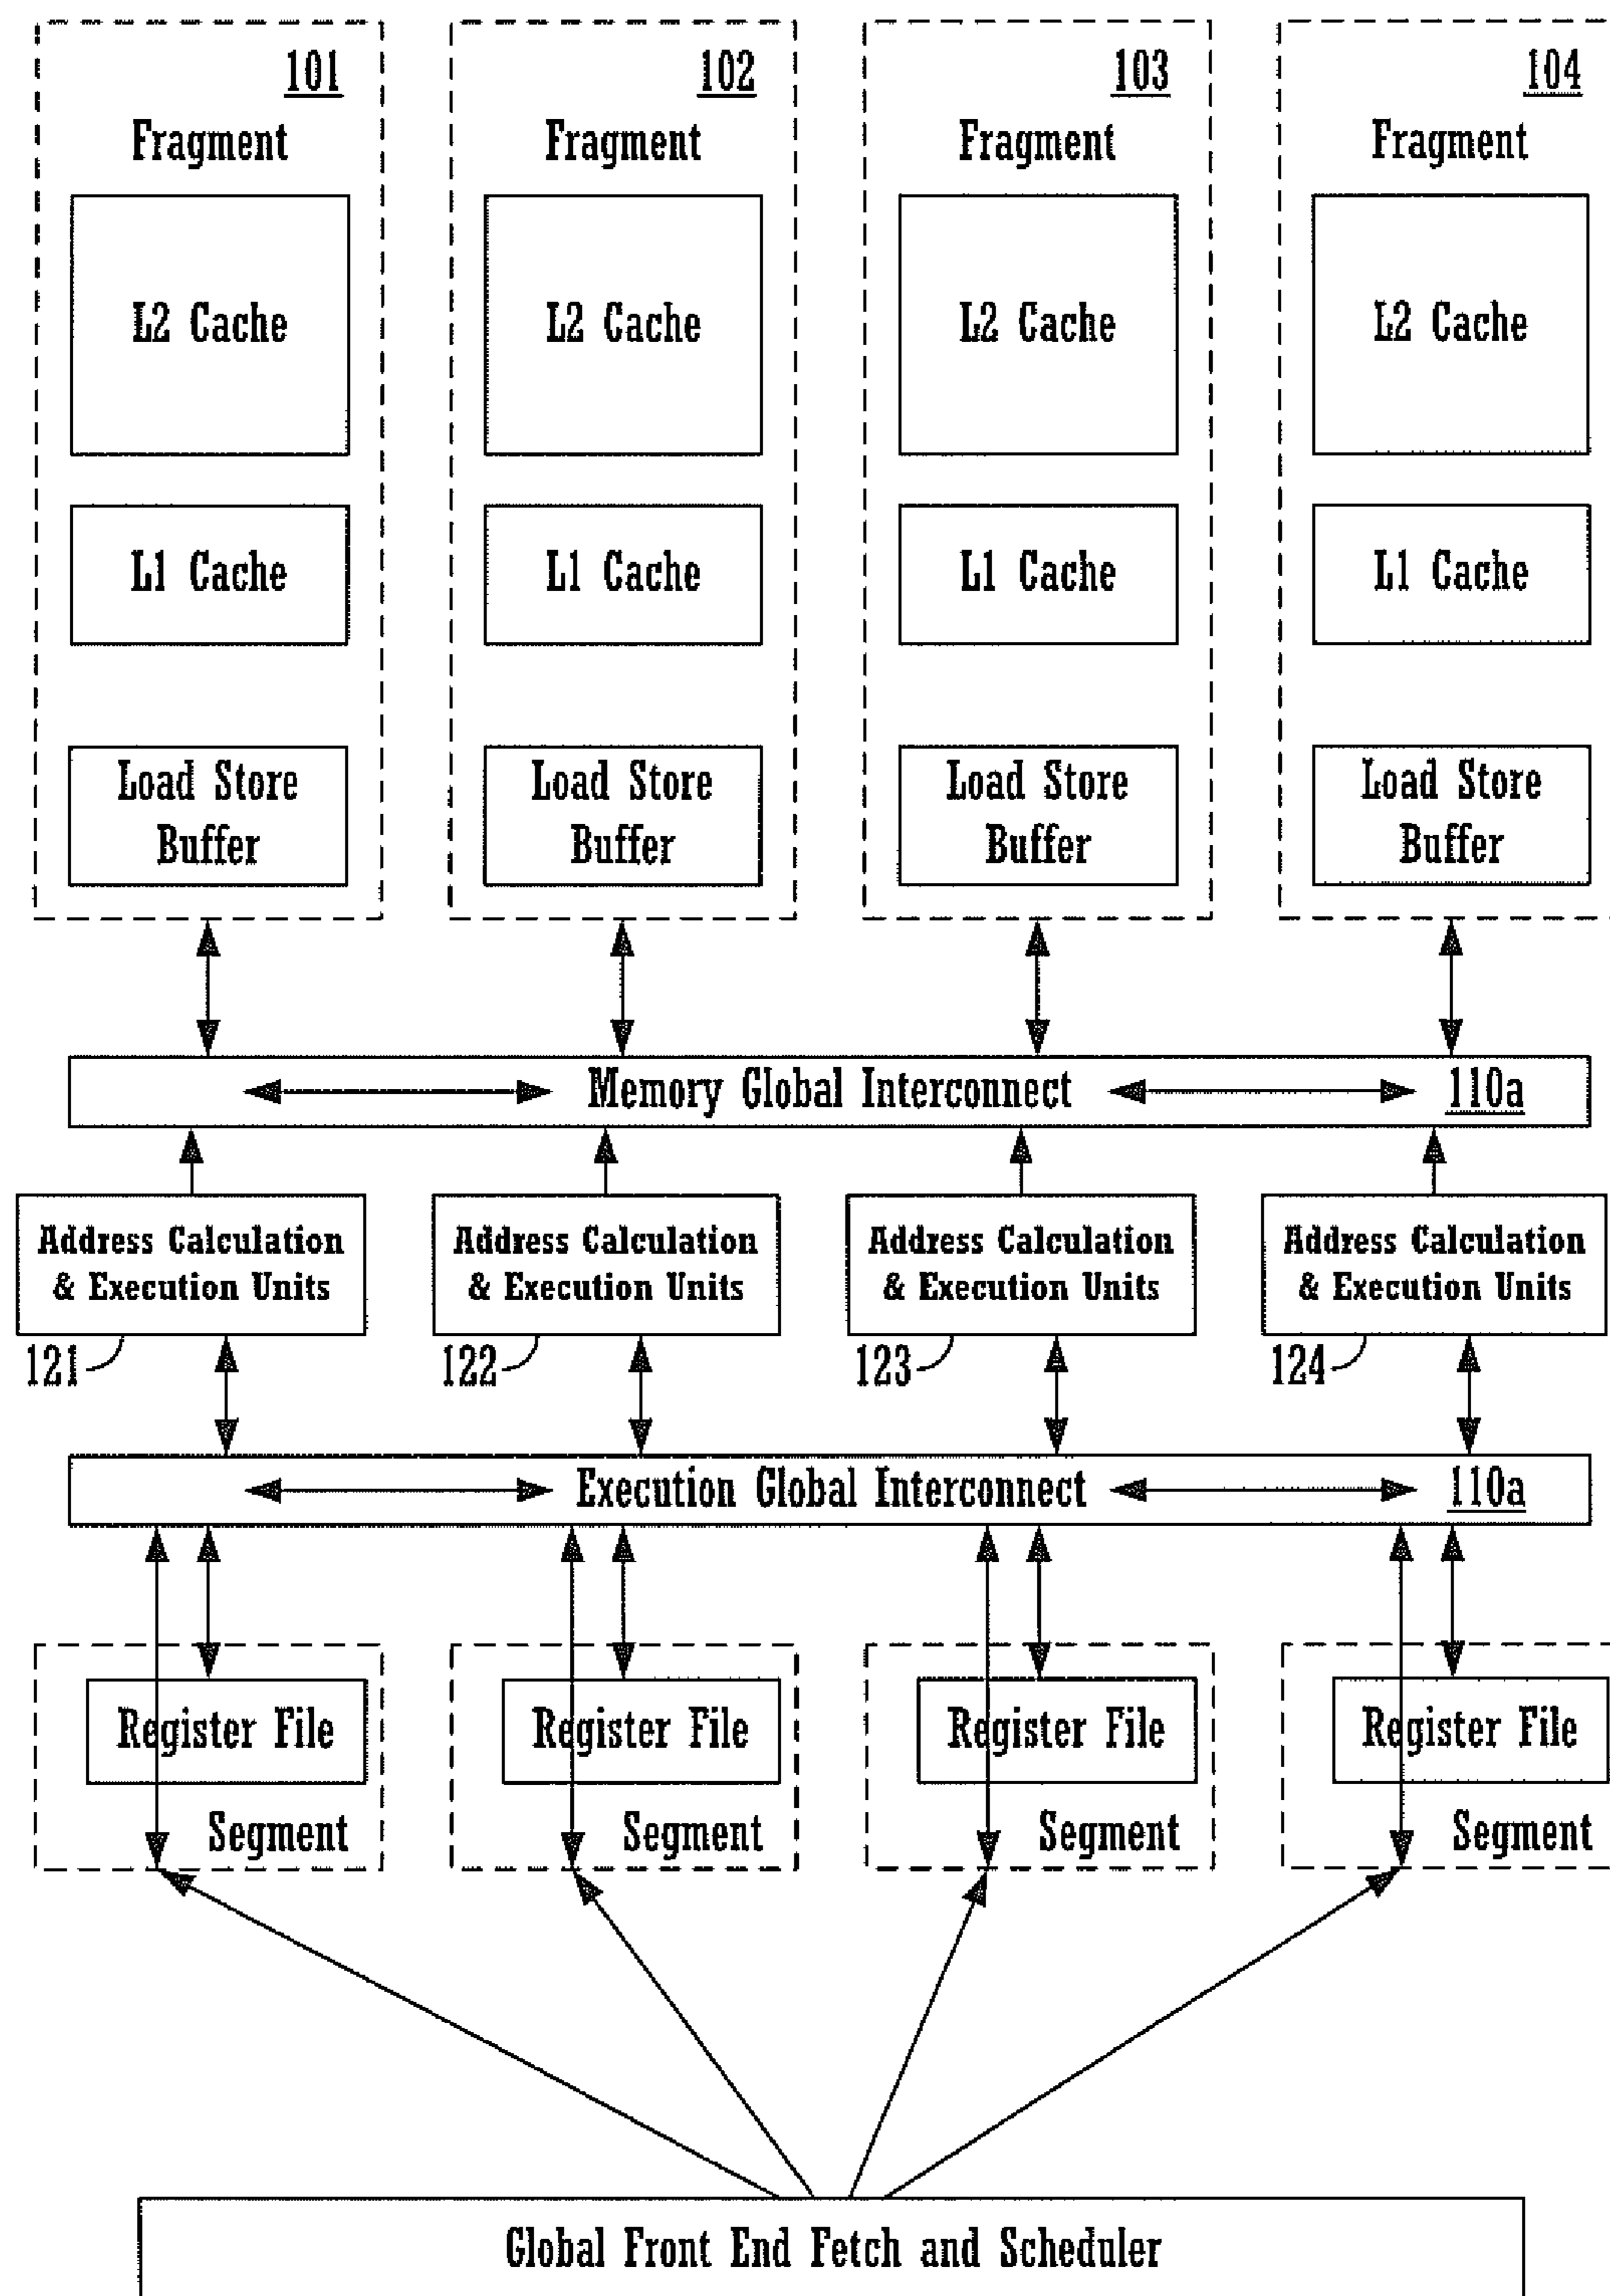


FIGURE 21

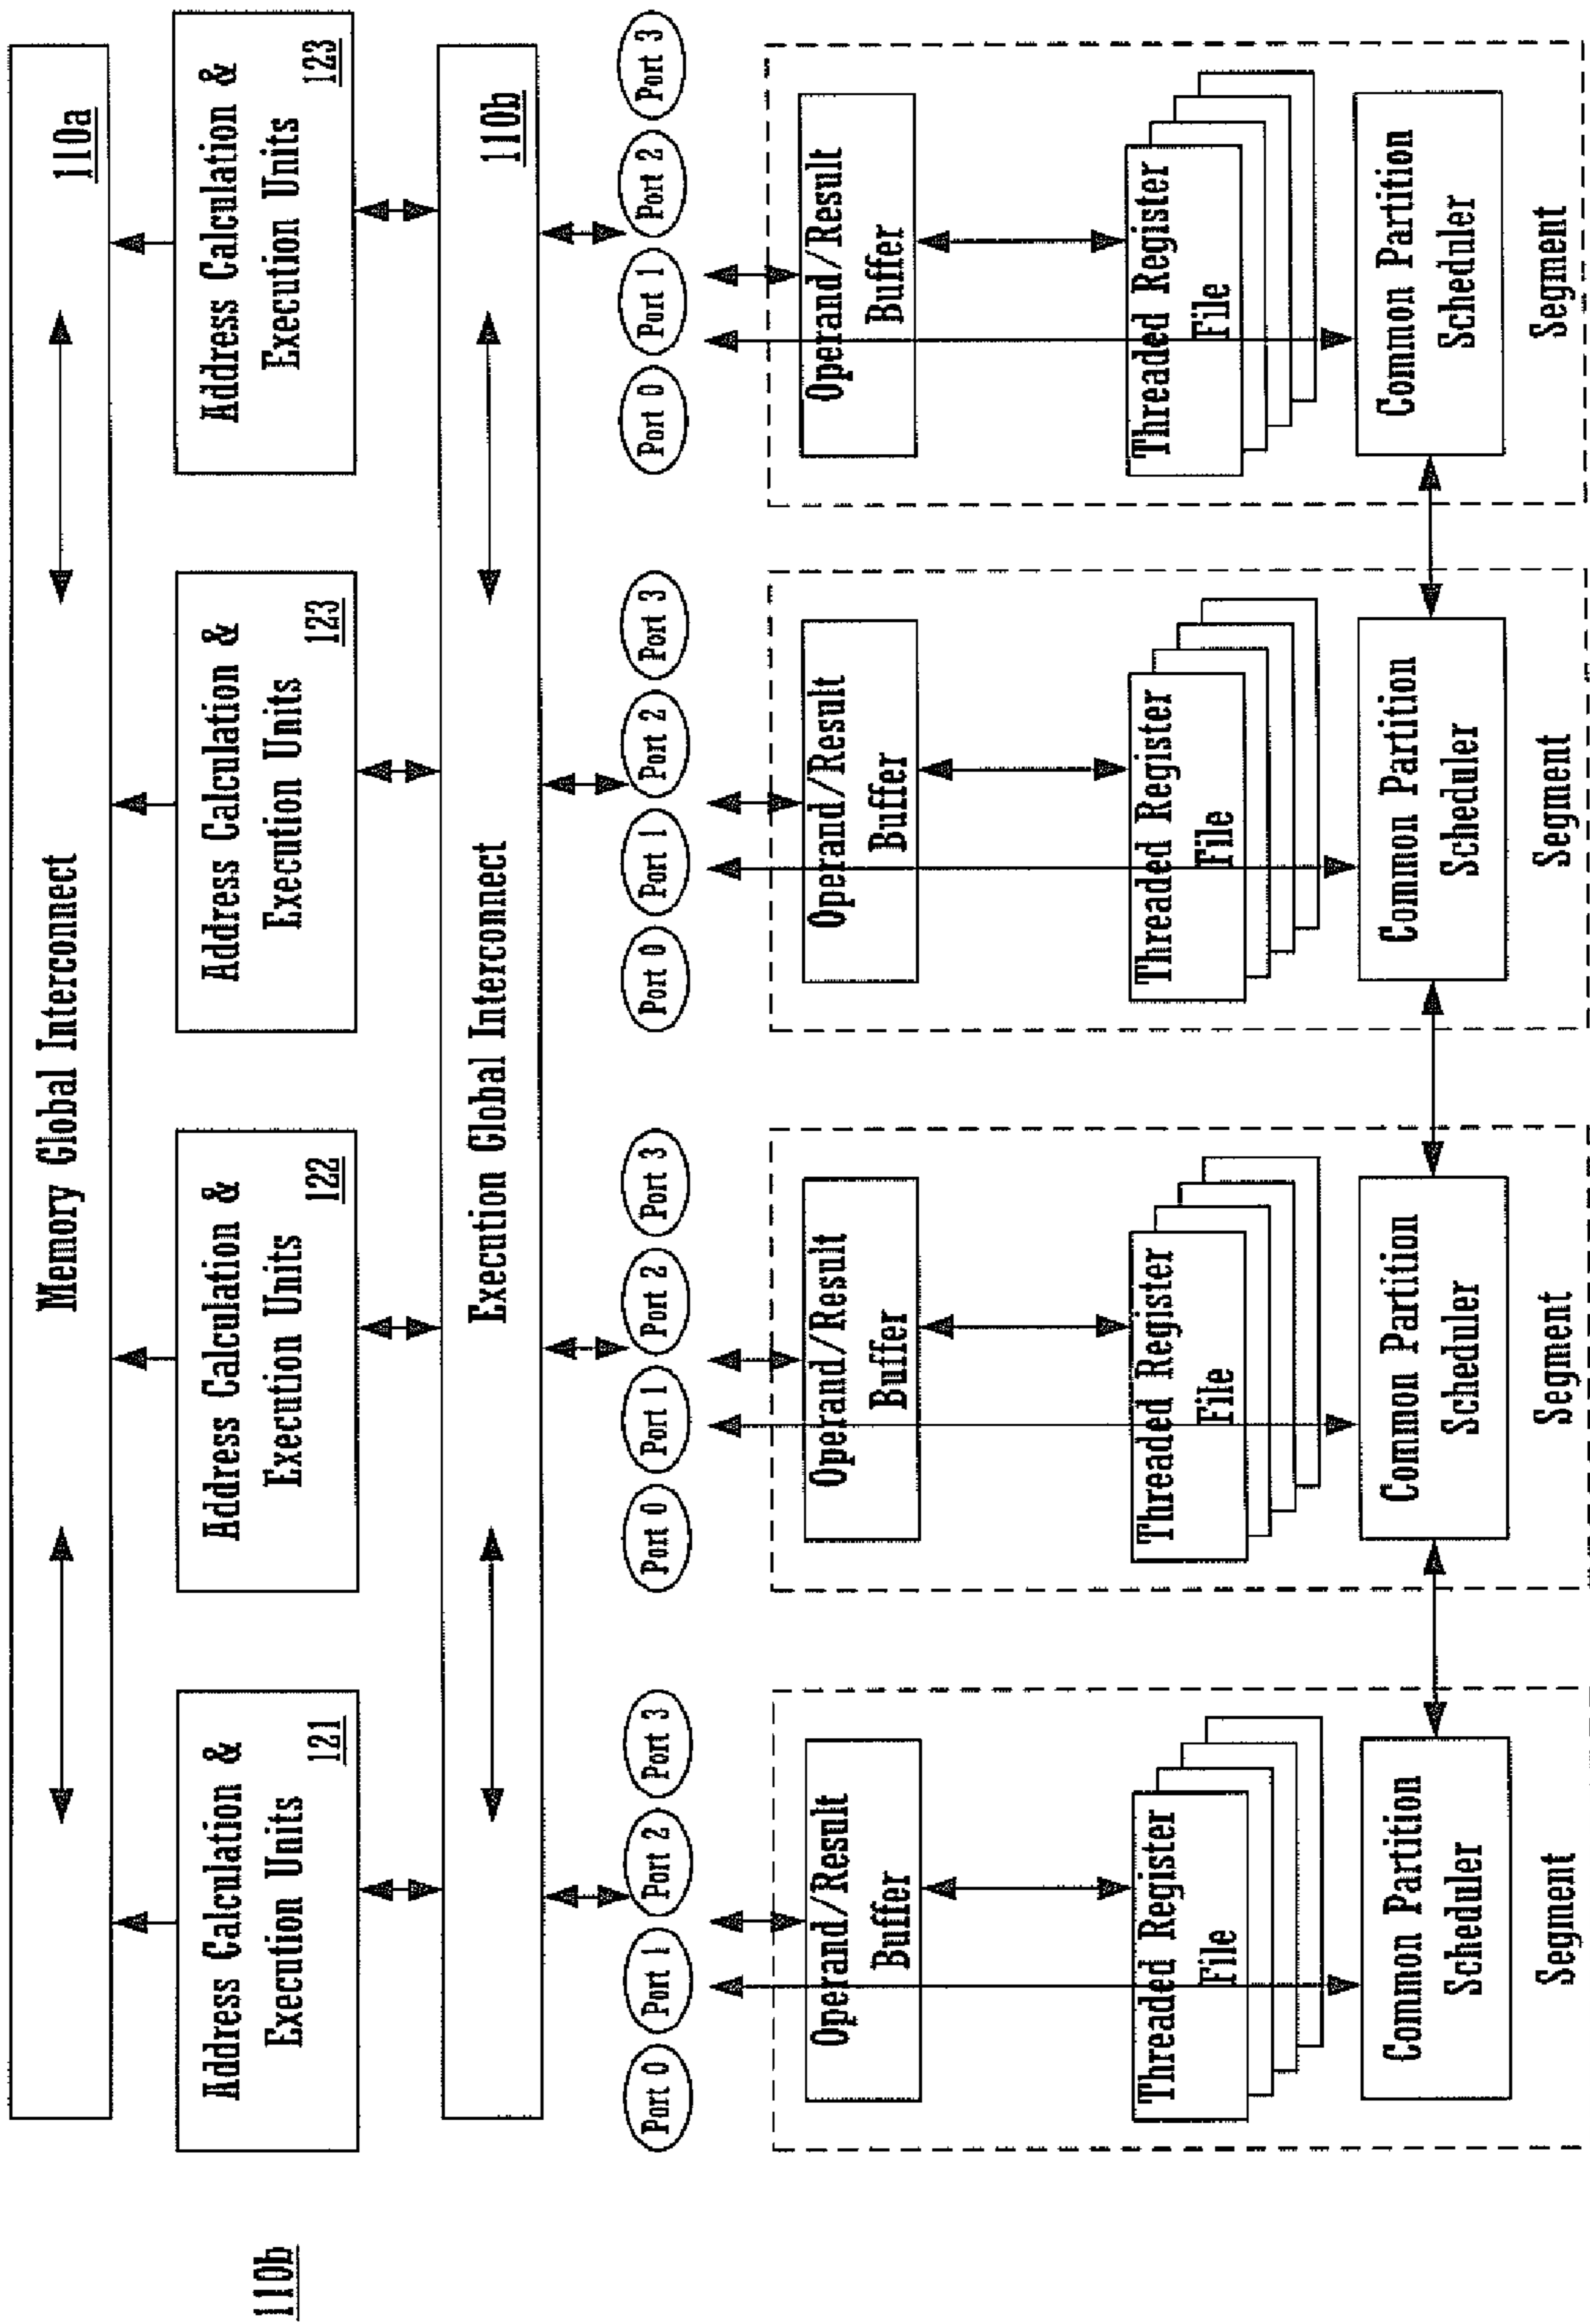


FIGURE 22

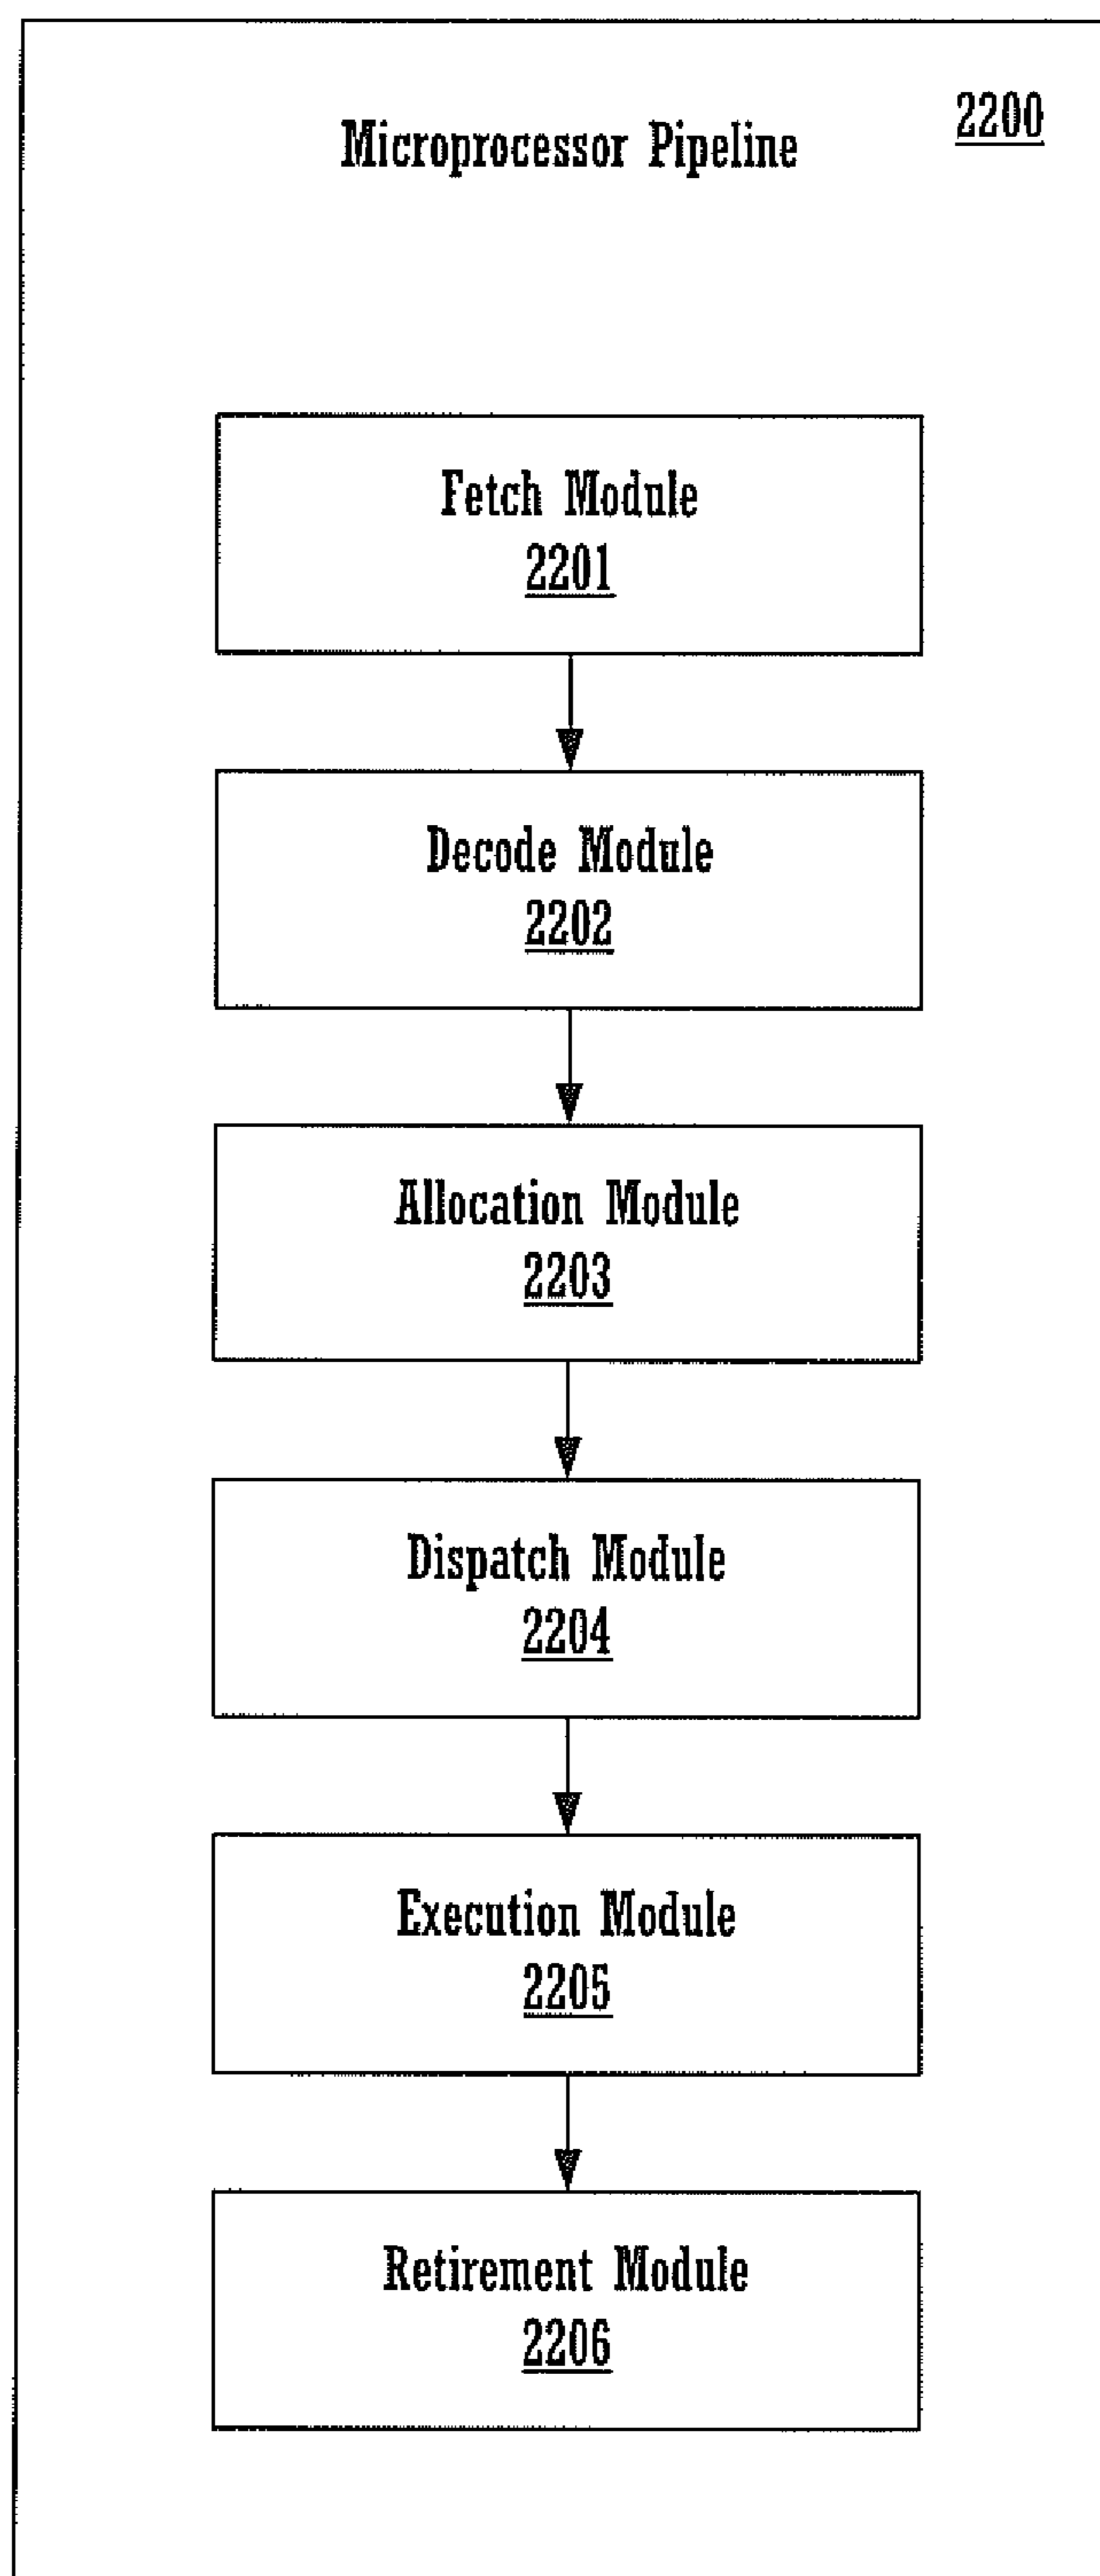


FIGURE 23

MICROPROCESSOR ACCELERATED CODE OPTIMIZER

CROSS REFERENCE TO RELATED APPLICATION

[0001] This application is related to co-pending commonly assigned US Patent Application serial number 2010/0161948, titled “APPARATUS AND METHOD FOR PROCESSING COMPLEX INSTRUCTION FORMATS IN A MULTITHREADED ARCHITECTURE SUPPORTING VARIOUS CONTEXT SWITCH MODES AND VIRTUALIZATION SCHEMES” by Mohammad A. Abdallah, filed on Jan. 5, 2010, and which is incorporated herein in its entirety.

[0002] This application is related to co-pending commonly assigned US Patent Application serial number 2009/0113170, titled “APPARATUS AND METHOD FOR PROCESSING AN INSTRUCTION MATRIX SPECIFYING PARALLEL IN DEPENDENT OPERATIONS” by Mohammad A. Abdallah, filed on Dec. 19, 2008, and which is incorporated herein in its entirety.

[0003] This application is related to co-pending commonly assigned U.S. Patent Application Ser. No. 61/384,198, titled “SINGLE CYCLE MULTI-BRANCH PREDICTION INCLUDING SHADOW CACHE FOR EARLY FAR BRANCH PREDICTION” by Mohammad A. Abdallah, filed on Sep. 17, 2010, and which is incorporated herein in its entirety.

[0004] This application is related to co-pending commonly assigned U.S. Patent Application Ser. No. 61/467,944, titled “EXECUTING INSTRUCTION SEQUENCE CODE BLOCKS BY USING VIRTUAL CORES INSTANTIATED BY PARTITIONABLE ENGINES” by Mohammad A. Abdallah, filed on Mar. 25, 2011, and which is incorporated herein in its entirety.

FIELD OF THE INVENTION

[0005] The present invention is generally related to digital computer systems, more particularly, to a system and method for selecting instructions comprising an instruction sequence.

BACKGROUND OF THE INVENTION

[0006] Processors are required to handle multiple tasks that are either dependent or totally independent. The internal state of such processors usually consists of registers that might hold different values at each particular instant of program execution. At each instant of program execution, the internal state image is called the architecture state of the processor.

[0007] When code execution is switched to run another function (e.g., another thread, process or program), then the state of the machine/processor has to be saved so that the new function can utilize the internal registers to build its new state. Once the new function is terminated then its state can be discarded and the state of the previous context will be restored and execution resumes. Such a switch process is called a context switch and usually includes 10's or hundreds of cycles especially with modern architectures that employ large number of registers (e.g., 64, 128, 256) and/or out of order execution.

[0008] In thread-aware hardware architectures, it is normal for the hardware to support multiple context states for a limited number of hardware-supported threads. In this case, the hardware duplicates all architecture state elements for each supported thread. This eliminates the need for context

switch when executing a new thread. However, this still has multiple draw backs, namely the area, power and complexity of duplicating all architecture state elements (i.e., registers) for each additional thread supported in hardware. In addition, if the number of software threads exceeds the number of explicitly supported hardware threads, then the context switch must still be performed.

[0009] This becomes common as parallelism is needed on a fine granularity basis requiring a large number of threads. The hardware thread-aware architectures with duplicate context-state hardware storage do not help non-threaded software code and only reduces the number of context switches for software that is threaded. However, those threads are usually constructed for coarse grain parallelism, and result in heavy software overhead for initiating and synchronizing, leaving fine grain parallelism, such as function calls and loops parallel execution, without efficient threading initiations/auto generation. Such described overheads are accompanied with the difficulty of auto parallelization of such codes using state of the art compiler or user parallelization techniques for non-explicitly/easily parallelized/threaded software codes.

SUMMARY OF THE INVENTION

[0010] In one embodiment the present invention is implemented as a method for accelerating code optimization in a microprocessor. The method includes fetching an incoming macroinstruction sequence using an instruction fetch component and transferring the fetched macroinstructions to a decoding component for decoding into microinstructions. Optimization processing is performed by reordering the microinstruction sequence into an optimized microinstruction sequence comprising a plurality of dependent code groups. The optimized microinstruction sequence is output to a microprocessor pipeline for execution. A copy of the optimized microinstruction sequence is stored into a sequence cache for subsequent use upon a subsequent hit to the optimized microinstruction sequence.

[0011] The foregoing is a summary and thus contains, by necessity, simplifications, generalizations and omissions of detail; consequently, those skilled in the art will appreciate that the summary is illustrative only and is not intended to be in any way limiting. Other aspects, inventive features, and advantages of the present invention, as defined solely by the claims, will become apparent in the non-limiting detailed description set forth below.

BRIEF DESCRIPTION OF THE DRAWINGS

[0012] The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements.

[0013] FIG. 1 shows an overview diagram of an allocation/issue stage of a microprocessor in accordance with one embodiment of the present invention.

[0014] FIG. 2 shows an overview diagram illustrating an optimization process in accordance with one embodiment of the present invention.

[0015] FIG. 3 shows a multistep optimization process in accordance with one embodiment of the present invention.

[0016] FIG. 4 shows a multistep optimization and instruction moving process in accordance with one embodiment of the present invention.

[0017] FIG. 5 shows a flowchart of the steps of an exemplary hardware optimization process in accordance with one embodiment of the present invention.

[0018] FIG. 6 shows a flowchart of the steps of an alternative exemplary hardware optimization process in accordance with one embodiment of the present invention.

[0019] FIG. 7 shows a diagram showing the operation of the CAM matching hardware and the priority encoding hardware of the allocation/issue stage in accordance with one embodiment of the present invention.

[0020] FIG. 8 shows a diagram illustrating optimized scheduling ahead of a branch in accordance with one embodiment of the present invention.

[0021] FIG. 9 shows a diagram illustrating optimized scheduling ahead of a store in accordance with one embodiment of the present invention.

[0022] FIG. 10 shows a diagram of an exemplary software optimization process in accordance with one embodiment of the present invention.

[0023] FIG. 11 shows a flow diagram of a SIMD software-based optimization process in accordance with one embodiment of the present invention.

[0024] FIG. 12 shows a flowchart of the operating steps of an exemplary SIMD software-based optimization process in accordance with one embodiment of the present invention.

[0025] FIG. 13 shows a software based dependency broadcast process in accordance with one embodiment of the present invention.

[0026] FIG. 14 shows an exemplary flow diagram that shows how the dependency grouping of instructions can be used to build variably bounded groups of dependent instructions in accordance with one embodiment of the present invention.

[0027] FIG. 15 shows a flow diagram depicting hierarchical scheduling of instructions in accordance with one embodiment of the present invention.

[0028] FIG. 16 shows a flow diagram depicting hierarchical scheduling of three slot dependency group instructions in accordance with one embodiment of the present invention.

[0029] FIG. 17 shows a flow diagram depicting hierarchical moving window scheduling of three slot dependency group instructions in accordance with one embodiment of the present invention.

[0030] FIG. 18 shows how the variably sized dependent chains (e.g., variably bounded groups) of instructions are allocated to a plurality of computing engines in accordance with one embodiment of the present invention.

[0031] FIG. 19 shows a flow diagram depicting block allocation to the scheduling queues and the hierarchical moving window scheduling of three slot dependency group instructions in accordance with one embodiment of the present invention.

[0032] FIG. 20 shows how the dependent code blocks (e.g., dependency groups or dependency chains) are executed on the engines in accordance with one embodiment of the present invention.

[0033] FIG. 21 shows an overview diagram of a plurality of engines and their components, including a global front end fetch & scheduler and register files, global interconnects and a fragmented memory subsystem for a multicore processor in accordance with one embodiment of the present invention.

[0034] FIG. 22 shows a plurality of segments, a plurality of segmented common partition schedulers and the interconnect

and the ports into the segments in accordance with one embodiment of the present invention.

[0035] FIG. 23 shows a diagram of an exemplary microprocessor pipeline in accordance with one embodiment of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

[0036] Although the present invention has been described in connection with one embodiment, the invention is not intended to be limited to the specific forms set forth herein. On the contrary, it is intended to cover such alternatives, modifications, and equivalents as can be reasonably included within the scope of the invention as defined by the appended claims.

[0037] In the following detailed description, numerous specific details such as specific method orders, structures, elements, and connections have been set forth. It is to be understood however that these and other specific details need not be utilized to practice embodiments of the present invention. In other circumstances, well-known structures, elements, or connections have been omitted, or have not been described in particular detail in order to avoid unnecessarily obscuring this description.

[0038] References within the specification to “one embodiment” or “an embodiment” are intended to indicate that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment of the present invention. The appearance of the phrase “in one embodiment” in various places within the specification are not necessarily all referring to the same embodiment, nor are separate or alternative embodiments mutually exclusive of other embodiments. Moreover, various features are described which may be exhibited by some embodiments and not by others. Similarly, various requirements are described which may be requirements for some embodiments but not other embodiments.

[0039] Some portions of the detailed descriptions, which follow, are presented in terms of procedures, steps, logic blocks, processing, and other symbolic representations of operations on data bits within a computer memory. These descriptions and representations are the means used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. A procedure, computer executed step, logic block, process, etc., is here, and generally, conceived to be a self-consistent sequence of steps or instructions leading to a desired result. The steps are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals of a computer readable storage medium and are capable of being stored, transferred, combined, compared, and otherwise manipulated in a computer system. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like.

[0040] It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the following discussions, it is appreciated that throughout the present invention, discussions utilizing terms such as “processing” or “accessing” or “writing” or “storing” or “replicating” or the like, refer to the action and processes of a computer system, or similar electronic computing device

that manipulates and transforms data represented as physical (electronic) quantities within the computer system's registers and memories and other computer readable media into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

[0041] In one embodiment the present invention is implemented as a method for accelerating code optimization in a microprocessor. The method includes fetching an incoming microinstruction sequence using an instruction fetch component and transferring the fetched macroinstructions to a decoding component for decoding into microinstructions. Optimization processing is performed by reordering the microinstruction sequence into an optimized microinstruction sequence comprising a plurality of dependent code groups. The optimized microinstruction sequence is output to a microprocessor pipeline for execution. A copy of the optimized microinstruction sequence is stored into a sequence cache for subsequent use upon a subsequent hit to the optimized microinstruction sequence.

[0042] FIG. 1 shows an overview diagram of an allocation/issue stage of a microprocessor 100 in accordance with one embodiment of the present invention. As illustrated in FIG. 1, the microprocessor 100 includes a fetch component 101, a native decode component 102, and instruction scheduling and optimizing component 110 and the remaining pipeline 105 of the microprocessor.

[0043] In the FIG. 1 embodiment, macroinstructions are fetched by a fetch component 101 and decoded into native microinstructions by the native decode component 102, which then provides the microinstructions to a microinstruction cache 121 and the instruction scheduling and optimizer component 110. In one embodiment, the fetched macroinstructions comprise a sequence of instructions that is assembled by predicting certain branches.

[0044] The macroinstruction sequence is decoded into a resulting microinstruction sequence by the native decode component 102. This microinstruction sequence is then transmitted to the instruction scheduling and optimizing component 110 through a multiplexer 103. The instruction scheduling and optimizer component functions by performing optimization processing by, for example, reordering certain instructions of the microinstruction sequence for more efficient execution. This results in an optimized microinstruction sequence that is then transferred to the remaining pipeline 105 (e.g., the allocation, dispatch, execution, and retirement stages, etc.) through the multiplexer 104. The optimized microinstruction sequence results in a faster and more efficient execution of the instructions.

[0045] In one embodiment, the macroinstructions can be instructions from a high level instruction set architecture, while the microinstructions are low level machine instructions. In another embodiment, the macroinstructions can be guest instructions from a plurality of different instruction set architectures (e.g., CISC like, x86, RISC like, MIPS, SPARC, ARM, virtual like, JAVA, and the like), while the microinstructions are low level machine instructions or instructions of a different native instruction set architecture. Similarly, in one embodiment, the macroinstructions can be native instructions of an architecture, and the microinstructions can be native microinstructions of that same architecture that have been reordered and optimized. For example X86 macro instructions and X86 micro-coded microinstructions.

[0046] In one embodiment, to accelerate the execution performance of code that is frequently encountered (e.g., hot code), copies of frequently encountered microinstruction sequences are cached in the microinstruction cache 121 and copies of frequently encountered optimized microinstruction sequences are cached within the sequence cache 122. As code is fetched, decoded, optimized, and executed, certain optimized microinstruction sequences can be evicted or fetched in accordance with the size of the sequence cache through the depicted eviction and fill path 130. This eviction and fill path allows for transfers of optimized microinstruction sequences to and from the memory hierarchy of the microprocessor (e.g., L1 cache, L2 cache, a special cacheable memory range, or the like).

[0047] It should be noted that in one embodiment, the microinstruction cache 121 can be omitted. In such an embodiment, the acceleration of hot code is provided by the storing of optimized microinstruction sequences within the sequence cache 122. For example, the space saved by omitting microinstruction cache 121 can be used to implement a larger sequence cache 122, for example.

[0048] FIG. 2 shows an overview diagram illustrating an optimization process in accordance with one embodiment of the present invention. The left-hand side of FIG. 2 shows an incoming microinstruction sequence as received from, for example, the native decode component 102 or the microinstruction cache 121. Upon first receiving these instructions, they are not optimized.

[0049] One objective of the optimization process is to locate and identify instructions that depend upon one another and move them into their respective dependency groups so that they can execute more efficiently. In one embodiment, groups of dependent instructions can be dispatched together so that they can execute more efficiently since their respective sources and destinations are grouped together for locality. It should be noted that this optimization processing can be used in both an out of order processor as well as an in order processor. For example, within an in order processor, instructions are dispatched in-order. However, they can be moved around so that dependent instructions are placed in respective groups so that groups can then execute independently, as described above.

[0050] For example, the incoming instructions include loads, operations and stores. For example, instruction 1 comprises an operation where source registers (e.g., register 9 and register 9) are added and the result stored in register 5. Hence, register 5 is a destination and register 9 and register 5 are sources. In this manner, the sequence of 16 instructions includes destination registers and source registers, as shown.

[0051] The FIG. 2 embodiment implements the reordering of instructions to create dependency groups where instructions that belong to a group are dependent upon one another. To accomplish this, an algorithm is executed that performs hazard checks with respect to the loads and stores of the 16 incoming instructions. For example, stores cannot move past earlier loads without dependency checks. Stores cannot pass earlier stores. Loads cannot pass earlier stores without dependency checks. Loads can pass loads. Instructions can pass prior path predicted branches (e.g., dynamically constructed branches) by using a renaming technique. In the case of non-dynamically predicted branches, movements of instructions need to consider the scopes of the branches. Each of the above rules can be implemented by adding virtual depen-

dependency (e.g., by artificially adding virtual sources or destinations to instructions to enforce the rules).

[0052] Referring still to FIG. 2, as described above, an objective of the optimization process is to locate dependent instructions and move them into a common dependency group. This process must be done in accordance with the hazard checking algorithm. The optimization algorithm is looking for instruction dependencies. The instruction dependencies further comprise true dependencies, output dependencies and anti-dependencies.

[0053] The algorithm begins by looking for true dependencies first. To identify true dependencies, each destination of the 16 instruction sequence is compared against other subsequent sources which occur later in the 16 instruction sequence. The subsequent instructions that are truly dependent on an earlier instruction are marked “_1” to signify their true dependence. This is shown in FIG. 2 by the instruction numbers that proceed from left to right over the 16 instruction sequence. For example, considering instruction number 4, the destination register R3 is compared against the subsequent instructions’ sources, and each subsequent source is marked “_1” to indicate that instruction’s true dependence. In this case, instruction 6, instruction 7, instruction 11, and instruction 15 are marked “_1”.

[0054] The algorithm then looks for output dependencies. To identify output dependencies, each destination is compared against other subsequent instructions’ destinations. And for each of the 16 instructions, each subsequent destination that matches is marked “1_” (e.g., sometimes referred to as a red one).

[0055] The algorithm then looks for anti-dependencies. To identify anti-dependencies, for each of the 16 instructions, each source is compared with earlier instructions’ sources to identify matches. If a match occurs, the instruction under consideration marks its self “1_” (e.g., sometimes referred to as a red one).

[0056] In this manner, the algorithm populates a dependency matrix of rows and columns for the sequence of 16 instructions. The dependency matrix comprises the marks that indicate the different types of dependencies for each of the 16 instructions. In one embodiment, the dependency matrix is populated in one cycle by using CAM matching hardware and the appropriate broadcasting logic. For example, destinations are broadcasted downward through the remaining instructions to be compared with subsequent instructions’ sources (e.g., true dependence) and subsequent instructions’ destinations (e.g., output dependence), while destinations can be broadcasted upward through the previous instructions to be compared with prior instructions’ sources (e.g., anti dependence).

[0057] The optimization algorithm uses the dependency matrix to choose which instructions to move together into common dependency groups. It is desired that instructions which are truly dependent upon one another be moved to the same group. Register renaming is used to eliminate anti-dependencies to allow those anti-dependent instructions to be moved. The moving is done in accordance with the above described rules and hazard checks. For example, stores cannot move past earlier loads without dependency checks. Stores cannot pass earlier stores. Loads cannot pass earlier stores without dependency checks. Loads can pass loads. Instructions can pass prior path predicted branches (e.g., dynamic the constructed branches) by using a renaming technique. In the case of non-dynamically predicted branches,

movements of instructions need to consider the scopes of the branches. Note for the description

[0058] In one embodiment, a priority encoder can be implemented to determine which instructions get moved to be grouped with other instructions. The priority encoder would function in accordance with the information provided by the dependency matrix.

[0059] FIG. 3 and FIG. 4 show a multistep optimization process in accordance with one embodiment of the present invention. In one embodiment, the optimization process is iterative, in that after instructions are moved in a first pass by moving their dependency column, the dependency matrix is repopulated and examined again for new opportunities to move instructions. In one embodiment, this dependency matrix population process is repeated three times. This is shown in FIG. 4, which show instructions that have been moved and then examined again looking for opportunities to move other instructions. The sequence of numbers on the right hand side of each of the 16 instructions shows the group that the instruction was in that it began the process with and the group that the instruction was in at the finish of the process, with the intervening group numbers in between. For example, FIG. 4 shows how instruction 6 was initially in group 4 but was moved to be in group 1.

[0060] In this manner, FIGS. 2 through 4 illustrate the operation of an optimization algorithm in accordance with one embodiment of the present invention. It should be noted that although FIGS. 2 through 4 illustrate an allocation/issue stage, this functionality can also be implemented in a local scheduler/dispatch stage.

[0061] FIG. 5 shows a flowchart of the steps of an exemplary hardware optimization process 500 in accordance with one embodiment of the present invention. As depicted in FIG. 5, the flowchart shows the operating steps of a optimization process as implemented in an allocation/issue stage of a microprocessor in accordance with one embodiment of the present invention.

[0062] Process 500 begins in step 501, where an incoming macroinstruction sequence is fetched using an instruction fetch component (e.g., fetch component 20 from FIG. 1). As described above, the fetched instructions comprise a sequence that is assembled by predicting certain instruction branches.

[0063] In step 502, the fetched macroinstructions are transferred to a decoding component for decoding into microinstructions. The macroinstruction sequence is decoded into a microinstruction sequence in accordance with the branch predictions. In one embodiment, the microinstruction sequence is then stored into a microinstruction cache.

[0064] In step 503, optimization processing is then conducted on the microinstruction sequence by reordering the microinstructions comprising sequence into dependency groups. The reordering is implemented by an instruction reordering component (e.g., the instruction scheduling and optimizer component 110). This process is described in the FIGS. 2 through 4.

[0065] In step 504, the optimized microinstruction sequence is an output to the microprocessor pipeline for execution. As described above, the optimized microinstruction sequence is forwarded to the rest of the machine for execution (e.g., remaining pipeline 105).

[0066] And subsequently, in step 505, a copy of the optimized microinstruction sequence is stored into a sequence cache for subsequent use upon a subsequent hit to that

sequence. In this manner, the sequence cache enables access to the optimized microinstruction sequences upon subsequent hits on those sequences, thereby accelerating hot code.

[0067] FIG. 6 shows a flowchart of the steps of an alternative exemplary hardware optimization process 600 in accordance with one embodiment of the present invention. As depicted in FIG. 6, the flowchart shows the operating steps of a optimization process as implemented in an allocation/issue stage of a microprocessor in accordance with an alternative embodiment of the present invention.

[0068] Process 600 begins in step 601, where an incoming macroinstruction sequence is fetched using an instruction fetch component (e.g., fetch component 20 from FIG. 1). As described above, the fetched instructions comprise a sequence that is assembled by predicting certain instruction branches.

[0069] In step 602, the fetched macroinstructions are transferred to a decoding component for decoding into microinstructions. The macroinstruction sequence is decoded into a microinstruction sequence in accordance with the branch predictions. In one embodiment, the microinstruction sequence is then stored into a microinstruction cache.

[0070] In step 603, the decoded micro instructions are stored into sequences in a micro instruction sequence cache. Sequences in the micro instruction cache are formed to start in accordance with basic block boundaries. These sequences are not optimized at this point.

[0071] In step 604, optimization processing is then conducted on the microinstruction sequence by reordering the microinstructions comprising sequence into dependency groups. The reordering is implemented by an instruction reordering component (e.g., the instruction scheduling and optimizer component 110). This process is described in the FIGS. 2 through 4.

[0072] In step 605, the optimized microinstruction sequence is an output to the microprocessor pipeline for execution. As described above, the optimized microinstruction sequence is forwarded to the rest of the machine for execution (e.g., remaining pipeline 105).

[0073] And subsequently, in step 606, a copy of the optimized microinstruction sequence is stored into a sequence cache for subsequent use upon a subsequent hit to that sequence. In this manner, the sequence cache enables access to the optimized microinstruction sequences upon subsequent hits on those sequences, thereby accelerating hot code.

[0074] FIG. 7 shows a diagram showing the operation of the CAM matching hardware and the priority encoding hardware of the allocation/issue stage in accordance with one embodiment of the present invention. As depicted in FIG. 7, destinations of the instructions are broadcast into the CAM array from the left. Three exemplary instruction destinations are shown. The lighter shaded CAMs (e.g. green) are for true dependency matches and output dependency matches, and thus the destinations are broadcast downward. The darker shaded CAMs (e.g. blue) anti-dependency matches, and thus the destinations are broadcast upward. These matches populate a dependency matrix, as described above. Priority encoders are shown on the right, and they function by scanning the row of CAMs to find the first match, either a “_1” or a “1_”. As described above in the discussions of FIGS. 2-4, the process can be implemented to be iterative. For example, if a “_1” is blocked by a “1_”, then that destination can be renamed and moved.

[0075] FIG. 8 shows a diagram illustrating optimized scheduling instructions ahead of a branch in accordance with one embodiment of the present invention. As illustrated in FIG. 8, a hardware optimized example is depicted alongside a traditional just-in-time compiler example. The left-hand side of FIG. 8 shows the original un-optimized code including the branch biased untaken, “Branch C to L1”. The middle column of FIG. 8 shows a traditional just-in-time compiler optimization, where registers are renamed and instructions are moved ahead of the branch. In this example, the just-in-time compiler inserts compensation code to account for those occasions where the branch biased decision is wrong (e.g., where the branch is actually taken as opposed to untaken). In contrast, the right column of FIG. 8 shows the hardware unrolled optimization. In this case, the registers are renamed and instructions are moved ahead of the branch. However, it should be noted that no compensation code is inserted. The hardware keeps track of whether branch biased decision is true or not. In case of wrongly predicted branches, the hardware automatically rolls back it’s state in order to execute the correct instruction sequence. The hardware optimizer solution is able to avoid the use of compensation code because in those cases where the branch is miss predicted, the hardware jumps to the original code in memory and executes the correct sequence from there, while flushing the miss predicted instruction sequence.

[0076] FIG. 9 shows a diagram illustrating optimized scheduling a load ahead of a store in accordance with one embodiment of the present invention. As illustrated in FIG. 9, a hardware optimized example is depicted alongside a traditional just-in-time compiler example. The left-hand side of FIG. 9 shows the original un-optimized code including the store, “R3←LD [R5]”. The middle column of FIG. 9 shows a traditional just-in-time compiler optimization, where registers are renamed and the load is moved ahead of the store. In this example, the just-in-time compiler inserts compensation code to account for those occasions where the address of the load instruction aliases the address of the store instruction (e.g., where the load movement ahead of the store is not appropriate). In contrast, the right column of FIG. 9 shows the hardware unrolled optimization. In this case, the registers are renamed and the load is also moved ahead of the store. However, it should be noted that no compensation code is inserted. In a case where moving the load ahead of the store is wrong, the hardware automatically rolls back it’s state in order to execute the correct instruction sequence. The hardware optimizer solution is able to avoid the use of compensation code because in those cases where the address alias-check branch is miss predicted, the hardware jumps to the original code in memory and executes the correct sequence from there, while flushing the miss predicted instruction sequence. In this case, the sequence assumes no aliasing. It should be noted that in one embodiment, the functionality diagrammed in FIG. 9 can be implemented by instruction scheduling and optimizer component 110 of FIG. 1. Similarly, it should be noted that in one embodiment, the functionality diagrammed in FIG. 9 can be implemented by the software optimizer 1000 described in FIG. 10 below.

[0077] Additionally, with respect to dynamically unrolled sequences, it should be noted that instructions can pass prior path predicted branches (e.g., dynamically constructed branches) by using renaming. In the case of non-dynamically predicted branches, movements of instructions should consider the scopes of the branches. Loops can be unrolled to the

extent desired and optimizations can be applied across the whole sequence. For example, this can be implemented by renaming destination registers of instructions moving across branches. One of the benefits of this feature is the fact that no compensation code or extensive analysis of the scopes of the branches is needed. This feature thus greatly speeds up and simplifies the optimization process.

[0078] Additional information concerning branch prediction and the assembling of instruction sequences can be found in commonly assigned U.S. Patent Application Ser. No. 61/384,198, titled "SINGLE CYCLE MULTI-BRANCH PREDICTION INCLUDING SHADOW CACHE FOR EARLY FAR BRANCH PREDICTION" by Mohammad A. Abdallah, filed on Sep. 17, 2010, which is incorporated herein in its entirety.

[0079] FIG. 10 shows a diagram of an exemplary software optimization process in accordance with one embodiment of the present invention. In the FIG. 10 embodiment, the instruction scheduling and optimizer component (e.g., component 110 of FIG. 1) is replaced by a software-based optimizer 1000.

[0080] In the FIG. 10 embodiment, the software optimizer 1000 performs the optimization processing that was performed by the hardware-based instruction scheduling and optimizer component 110. The software optimizer maintains a copy of optimized sequences in the memory hierarchy (e.g., L1, L2, system memory). This allows the software optimizer to maintain a much larger collection of optimized sequences in comparison to what is stored in the sequence cache.

[0081] It should be noted that the software optimizer 1000 can comprise code residing in the memory hierarchy as both input to the optimization and output from the optimization process.

[0082] It should be noted that in one embodiment, the microinstruction cache can be omitted. In such an embodiment, only the optimized microinstruction sequences are cached.

[0083] FIG. 11 shows a flow diagram of a SIMD software-based optimization process in accordance with one embodiment of the present invention. The top of FIG. 11 shows how the software-based optimizer examines each instruction of an input instruction sequence. FIG. 11 shows how a SIMD compare can be used to match one to many (e.g., SIMD byte compare a first source "Src1" to all second source bytes "Src2"). In one embodiment, Src1 contains the destination register of any instruction and Src2 contains one source from each other subsequent instruction. Matching is done for every destination with all subsequent instruction sources (e.g., true dependence checking). This is a pairing match that indicates a desired group for the instruction. Matching is done between each destination and every subsequent instruction destination (e.g., output dependence checking). This is a blocking match that can be resolved with renaming. Matching is done between each destination and every prior instruction source (e.g., anti dependence checking). This is a blocking match that can be resolved by renaming. The results are used to populate the rows and columns of the dependency matrix.

[0084] FIG. 12 shows a flowchart of the operating steps of an exemplary SIMD software-based optimization process 1200 in accordance with one embodiment of the present invention. Process 1200 is described in the context of the flow diagram of FIG. 9.

[0085] In step 1201, an input sequence of instructions is accessed by using a software-based optimizer instantiated memory.

[0086] In step 1202, a dependency matrix is populated, using SIMD instructions, with dependency information extracted from the input sequence of instructions by using a sequence of SIMD compare instructions.

[0087] In step 1203, the rows of the matrix are scanned from right to left for the first match (e.g., dependency mark).

[0088] In step 1204, each of the first matches are analyzed to determine the type of the match.

[0089] In step 1205, if the first marked match is a blocking dependency, renaming is done for this destination.

[0090] In step 1206, all first matches for each row of the matrix are identified and the corresponding column for that match is moved to the given dependency group.

[0091] In step 1207, the scanning process is repeated several times to reorder instructions comprising the input sequence to produce an optimized output sequence.

[0092] In step 1208, the optimized instruction sequence is output to the execution pipeline of the microprocessor for execution.

[0093] In step 1209, the optimized output sequence is stored in a sequence cache for subsequent consumption (e.g., to accelerate hot code).

[0094] It should be noted that the software optimization can be done serially with the use of SIMD instructions. For example, the optimization can be implemented by processing one instruction at a time scanning instructions' sources and destinations (e.g., from earlier instructions to subsequent instructions in a sequence). The software uses SIMD instructions to compare in parallel current instruction sources and destinations with prior instruction sources and destinations in accordance with the above described optimization algorithm and SIMD instructions (e.g. to detect true dependencies, output dependencies and anti-dependencies).

[0095] FIG. 13 shows a software based dependency broadcast process in accordance with one embodiment of the present invention. The FIG. 13 embodiment shows a flow diagram of an exemplary software scheduling process that processes groups of instructions without the expense of a full parallel hardware implementation as described above. However, the FIG. 13 embodiment can still use SIMD to process smaller groups of instructions in parallel.

[0096] The software scheduling process of FIG. 13 proceeds as follows. First, the process initializes three registers. The process takes instruction numbers and loads them into a first register. The process then takes destination register numbers and loads them into a second register. The process then takes the values in the first register and broadcasts them to a position in the third result register in accordance with a position number in the second register. The process then overwrites, going from left to right in the second register, the leftmost value will overwrite a right value in those instances where broadcast goes to the same position in the result register. Positions in the third register that have not been written to are bypassed. This information is used to populate a dependency matrix.

[0097] The FIG. 13 embodiment also shows the manner in which an input sequence of instructions can be processed as a plurality of groups. For example, a 16 instruction input sequence can be processed as a first group of 8 instructions and a second group of 8 instructions. With the first group, instruction numbers are loaded into the first register, instruc-

tion destination numbers are loaded into the second register, and the values in the first register are broadcast to positions in the third register (e.g., the result register) in accordance with the position number in the second register (e.g., a group broadcast). Positions in the third register that have not been written to are bypassed. The third register now becomes a base for the processing of the second group. For example, the result register from group 1 now becomes the result register for the processing of group two.

[0098] With the second group, instruction numbers are loaded into the first register, instruction destination numbers are loaded into the second register, and the values in the first register are broadcast to positions in the third register (e.g., the result register) in accordance with the position number in the second register. Positions in the third register can overwrite the result that was written during the processing of the first group. Positions in the third register that have not been written to are bypassed. In this manner, the second group updates the base from the first group, and thereby produces a new base for the processing of a third group, and so on.

[0099] Instructions in the second group can inherit dependency information generated in the processing of the first group. It should be noted that the entire second group does not have to be processed to update dependency in the result register. For example, dependency for instruction 12 can be generated in the processing of the first group, and then processing instructions in the second group up to instruction 11. This updates the result register to a state up to instruction 12. In one embodiment, a mask can be used to prevent the updates for the remaining instructions of the second group (e.g., instructions 12 through 16). To determine dependency for instruction 12, the result register is examined for R2 and R5. R5 will be updated with instruction 1, and R2 will be updated with instruction 11. It should be noted that in a case where all of group 2 is processed, R2 will be updated with instruction 15.

[0100] Additionally, it should be noted that all the instructions of the second group (e.g., instructions 9-16) can be processed independent of one another. In such case, the instructions of the second group depend only on the result register of the first group. The instructions of the second group can be processed in parallel once the result register is updated from the processing of the first group. In this manner, groups of instructions can be processed in parallel, one after another. In one embodiment, each group is processed using a SIMD instruction (e.g., a SIMD broadcast instruction), thereby processing all instructions of said each group in parallel.

[0101] FIG. 14 shows an exemplary flow diagram that shows how the dependency grouping of instructions can be used to build variably bounded groups of dependent instructions in accordance with one embodiment of the present invention. In the descriptions of FIGS. 2 through 4, the group sizes were constrained, in those cases three instructions per group. FIG. 14 shows how instructions can be reordered into variably sized groups, which then can be allocated to a plurality of computing engines. For example, FIG. 14 shows 4 engines. Since the groups can be variably sized depending on their characteristics, engine 1 can be allocated a larger group than, for example, engine 2. This can occur, for example, in a case where engine 2 has an instruction that is not particularly dependent upon the other instructions in that group.

[0102] FIG. 15 shows a flow diagram depicting hierarchical scheduling of instructions in accordance with one embodi-

ment of the present invention. As described above, dependency grouping of instructions can be used to build variably bounded groups. FIG. 15 shows the feature wherein various levels of dependency exist within a dependency group. For example, instruction 1 does not depend on any other instruction within this instruction sequence, therefore making instruction 1 an L0 dependency level. However, instruction 4 depends on instruction 1, therefore making instruction 4 an L1 dependency level. In this manner, each of the instructions of an instruction sequence is assigned a dependency level as shown.

[0103] The dependency level of each instruction is used by a second-level hierarchical scheduler to dispatch instructions in such a manner as to ensure resources are available for dependent instructions to execute. For example, in one embodiment, L0 instructions are loaded into instruction queues that are processed by the second-level schedulers 1-4. The L0 instructions are loaded such that they are in front of each of the queues, the L1 instructions are loaded such that they follow in each of the queues, L2 instructions follow them, and so on. This is shown by the dependency levels, from L0 to Ln in FIG. 15. The hierarchical scheduling of the schedulers 1-4 advantageously utilizes the locality-in-time and the instruction-to-instruction dependency to make scheduling decisions in an optimal way.

[0104] In this manner, embodiments of the present invention intimate dependency group slot allocation for the instructions of the instruction sequence. For example, to implement an out of order microarchitecture, the dispatching of the instructions of the instruction sequence is out of order. In one embodiment, on each cycle, instruction readiness is checked. An instruction is ready if all instructions that it depends upon have previously dispatched. A scheduler structure functions by checking those dependencies. In one embodiment, the scheduler is a unified scheduler and all dependency checking is performed in the unified scheduler structure. In another embodiment, the scheduler functionality is distributed across the dispatch queues of execution units of a plurality of engines. Hence, in one embodiment the scheduler is unified while in another embodiment the scheduler is distributed. With both of these solutions, each instruction source is checked against the dispatch instructions' destination every cycle.

[0105] Thus, FIG. 15 shows the hierarchical scheduling as performed by embodiments of the present invention. As described above, instructions are first grouped to form dependency chains (e.g., dependency groups). The formation of these dependency chains can be done statically or dynamically by software or hardware. Once these dependency chains have been formed, they can be distributed/dispatched to an engine. In this manner, grouping by dependency allows for out of order scheduling of in order formed groups. Grouping by dependency also distributes entire dependency groups onto a plurality of engines (e.g., cores or threads). Grouping by dependency also facilitates hierarchical scheduling as described above, where dependent instructions are grouped in a first step and then scheduled in a second step.

[0106] It should be noted that the functionality diagrammed in the FIGS. 14-19 can function independently from any method by which instructions are grouped (e.g., whether the grouping functionality is implemented in hardware, software, etc.). Additionally, the dependency groups shown in FIGS. 14-19 can comprise a matrix of independent groups, where each group further comprises dependent instructions. Addi-

tionally, it should be noted that the schedulers can also be engines. In such embodiment, each of the schedulers 1-4 can be incorporated within its respective engine (e.g., as shown in FIG. 22 where each segment includes a common partition scheduler).

[0107] FIG. 16 shows a flow diagram depicting hierarchical scheduling of three slot dependency group instructions in accordance with one embodiment of the present invention. As described above, dependency grouping of instructions can be used to build variably bounded groups. In this embodiment, the dependency groups comprise three slots. FIG. 16 shows the various levels of dependency even within a three slot dependency group. As described above, instruction 1 does not depend on any other instruction within this instruction sequence, therefore making instruction 1 an L0 dependency level. However, instruction 4 depends on instruction 1, therefore making instruction 4 an L1 dependency level. In this manner, each of the instructions of an instruction sequence is assigned a dependency level as shown.

[0108] As described above, the dependency level of each instruction is used by a second-level hierarchical scheduler to dispatch instructions in such a manner as to ensure resources are available for dependent instructions to execute. L0 instructions are loaded into instruction queues that are processed by the second-level schedulers 1-4. The L0 instructions are loaded such that they are in front of each of the queues, the L1 instructions are loaded such that they follow in each of the queues, L2 instructions follow them, and so on, as shown by the dependency levels, from L0 to Ln in FIG. 16. It should be noted that group number four (e.g., the fourth group from the top) begins at L2 even though it is a separate group. This is because instruction 7 depends from instruction 4, which depends from instruction 1, thereby giving instructions 7 an L2 dependency.

[0109] In this manner, FIG. 16 shows how every three dependent instructions are scheduled together on a given one of the schedulers 1-4. The second-level groups it scheduled behind the first level groups, then the groups are rotated.

[0110] FIG. 17 shows a flow diagram depicting hierarchical moving window scheduling of three slot dependency group instructions in accordance with one embodiment of the present invention. In this embodiment, the hierarchical scheduling for the three slot dependency groups is implemented via a unified moving window scheduler. A moving window scheduler processes the instructions in the queues to dispatch instructions in such a manner as to ensure resources are available for dependent instructions to execute. As described above, L0 instructions are loaded into instruction queues that are processed by the second-level schedulers 1-4. The L0 instructions are loaded such that they are in front of each of the queues, the L1 instructions are loaded such that they follow in each of the queues, L2 instructions follow them, and so on, as shown by the dependency levels, from L0 to Ln in FIG. 17. The moving window illustrates how L0 instructions can be dispatched from each of the queues even though they may be more in one queue than another. In this manner, the moving window scheduler dispatches instructions as the queues flow from left to right as illustrated in FIG. 17.

[0111] FIG. 18 shows how the variably sized dependent chains (e.g., variably bounded groups) of instructions are allocated to a plurality of computing engines in accordance with one embodiment of the present invention.

[0112] As depicted in FIG. 18, the processor includes an instruction scheduler component 10 and a plurality of engines

11-14. The instruction scheduler component generates code blocks and inheritance vectors to support the execution of dependent code block (e.g., variably bound group) on their respective engines. Each of the dependent code blocks can belong to the same logical core/thread or to different logical cores/threads. The instruction scheduler component will process the dependent code blocks to generate and respective inheritance vectors. These dependent code blocks and respective inheritance vectors are allocated to the particular engines 11-14 as shown. A global interconnect 30 supports a necessary communication across each of the engines 11-14. It should be noted that the functionality for the dependency grouping of instructions to build variably bounded groups of dependent instructions as described above in the discussion FIG. 14 is implemented by the instruction scheduler component 10 of the FIG. 18 embodiment.

[0113] FIG. 19 shows a flow diagram depicting block allocation to the scheduling queues and the hierarchical moving window scheduling of three slot dependency group instructions in accordance with one embodiment of the present invention. As described above, the hierarchical scheduling for the three slot dependency groups can be implemented via a unified moving window scheduler. FIG. 19 shows how dependency groups become blocks that are loaded into the scheduling queues. In FIG. 19 embodiment, two independent groups can be loaded in each queue as half blocks. This is shown at the top of FIG. 19 where group 1 forms one half block and group 4 forms another half block that is loaded into the first scheduling queue.

[0114] As described above, moving window scheduler processes the instructions in the queues to dispatch instructions in such a manner as to ensure resources are available for dependent instructions to execute. The bottom of FIG. 19 shows how L0 instructions are loaded into instruction queues that are processed by the second-level schedulers.

[0115] FIG. 20 shows how the dependent code blocks (e.g., dependency groups or dependency chains) are executed on the engines 11-14 in accordance with one embodiment of the present invention. As described above, instruction scheduler component generates code blocks and inheritance vectors to support the execution of dependent code blocks (e.g., variably bound group, three slot group, etc.) on their respective engines. As described above in FIG. 19, FIG. 20 further shows how two independent groups can be loaded into each engine as code blocks. FIG. 20 shows how these code blocks are dispatched to the engines 11-14, where the dependent instructions execute on the stacked (e.g., serially connected) execution units of each engine. For example, in the first dependency group, or code block, on the top left of FIG. 20, the instructions are dispatched to the engine 11 wherein they are stacked on the execution unit in order of their dependency such that L0 is stacked on top of L1 which is further stacked on L2. In so doing, the results of L0 to flow to the execution unit of L1 which can then flow to the execution of L2.

[0116] In this manner, the dependency groups shown in FIG. 20 can comprise a matrix of independent groups, where each group further comprises dependent instructions. The benefit of the groups being independent is the ability to dispatch and execute them in parallel and the attribute whereby the need for communication across the interconnect between the engines is minimized. Additionally, it should be noted that the execution units shown in the engines 11-14 can comprise a CPU or a GPU.

[0117] In accordance with embodiments of the present invention, it should be appreciated that instructions are abstracted into dependency groups or blocks or instruction matrices in accordance with their dependencies. Grouping instructions in accordance with their dependencies facilitates a more simplified scheduling process with a larger window of instructions (e.g., a larger input sequence of instructions). The grouping as described above removes the instruction variation and abstracts such variation uniformly, thereby allowing the implementation of simple, homogenous and uniform scheduling decision-making. The above described grouping functionality increases the throughput of the scheduler without increasing the complexity of the scheduler. For example, in a scheduler for four engines, the scheduler can dispatch four groups where each group has three instructions. In so doing, the scheduler only handles four lanes of super scalar complexity while dispatching 12 instructions. Furthermore, each block can contain parallel independent groups which further increase the number of dispatched instructions.

[0118] FIG. 21 shows an overview diagram of a plurality of engines and their components, including a global front end fetch & scheduler and register files, global interconnects and a fragmented memory subsystem for a multicore processor in accordance with one embodiment of the present invention. As depicted in FIG. 21, four memory fragments 101-104 are shown. The memory fragmentation hierarchy is the same across each cache hierarchy (e.g., L1 cache, L2 cache, and the load store buffer). Data can be exchanged between each of the L1 caches, each of the L2 caches and each of the load store buffers through the memory global interconnect 110a.

[0119] The memory global interconnect comprises a routing matrix that allows a plurality of cores (e.g., the address calculation and execution units 121-124) to access data that may be stored at any point in the fragmented cache hierarchy (e.g., L1 cache, load store buffer and L2 cache). FIG. 21 also depicts the manner whereby each of the fragments 101-104 can be accessed by address calculation and execution units 121-124 through the memory global interconnect 110a.

[0120] The execution global interconnect 110b similarly comprises a routing matrix allows the plurality of cores (e.g., the address calculation and execution units 121-124) to access data that may be stored at any of the segmented register files. Thus, the cores have access to data stored in any of the fragments and to data stored in any of the segments through the memory global interconnect 110a or the execution global interconnect 110b.

[0121] FIG. 21 further shows a global front end fetch & scheduler which has a view of the entire machine and which manages the utilization of the register files segments and the fragmented memory subsystem. Address generation comprises the basis for fragment definition. The global front end Fetch & scheduler functions by allocating instruction sequences to each segment.

[0122] FIG. 22 shows a plurality of segments, a plurality of segmented common partition schedulers and the interconnect and the ports into the segments in accordance with one embodiment of the present invention. As depicted in FIG. 22, each segment is shown with a common partition scheduler. The common partition scheduler functions by scheduling instructions within its respective segment. These instructions were in turn received from the global front end fetch and scheduler. In this embodiment, the common partition scheduler is configured to function in cooperation with the global front end fetch and scheduler. The segments are also shown

with 4 read write ports that provide read/write access to the operand/result buffer, threaded register file, and common partition or scheduler.

[0123] In one embodiment, a non-centralized access process is implemented for using the interconnects and the local interconnects employ the reservation adder and a threshold limiter control access to each contested resource, in this case, the ports into each segment. In such an embodiment, to access a resource, a core needs to reserve the necessary bus and reserve the necessary port.

[0124] FIG. 23 shows a diagram of an exemplary microprocessor pipeline 2300 in accordance with one embodiment of the present invention. The microprocessor pipeline 2300 includes a fetch module 2301 that implements the functionality of the process for identifying and extracting the instructions comprising an execution, as described above. In the FIG. 23 embodiment, the fetch module is followed by a decode module 2302, an allocation module 2303, a dispatch module 2304, an execution module 2305 and a retirement module 2306. It should be noted that the microprocessor pipeline 2300 is just one example of the pipeline that implements the functionality of embodiments of the present invention described above. One skilled in the art would recognize that other microprocessor pipelines can be implemented that include the functionality of the decode module described above.

[0125] For purposes of explanation, the foregoing description refers to specific embodiments that are not intended to be exhaustive or to limit the current invention. Many modifications and variations are possible consistent with the above teachings. Embodiments were chosen and described in order to best explain the principles of the invention and its practical applications, so as to enable others skilled in the art to best utilize the invention and its various embodiments with various modifications as may be suited to their particular uses.

What is claimed is:

1. In a microprocessor, a method for accelerating code optimization, comprising:

fetching an incoming microinstruction sequence using an instruction fetch component;

transferring the fetched macro instructions to a decoding component for decoding into microinstructions;

performing optimization processing by reordering the microinstruction sequence into an optimized microinstruction sequence comprising a plurality of dependent code groups;

outputting the optimized microinstruction sequence to a microprocessor pipeline for execution; and

storing a copy of the optimized microinstruction sequence into a sequence cache for subsequent use upon a subsequent hit optimized microinstruction sequence.

2. The method of claim 1, wherein a copy of the decoded microinstructions are stored in a microinstruction cache.

3. The method of claim 1, wherein the optimization processing is performed using an allocation and issue stage of the microprocessor.

4. The method of claim 3, wherein the allocation and issue stage further comprises an instruction scheduling and optimizer component that reorders the microinstruction sequence into the optimized micro instruction sequence.

5. The method of claim 1, wherein the optimization processing further comprises dynamically unrolling microinstruction sequences.

6. The method of claim 1, wherein the optimization processing is implemented through a plurality of iterations.

7. The method of claim 1, wherein the optimization processing is implemented through a register renaming process to enable the reordering.

8. A microprocessor, comprising:

an instruction fetch component for fetching an incoming microinstruction sequence;

a decoding component coupled to the instruction fetch component to receive the fetched macro instruction sequence and decode into a microinstruction sequence;

an allocation and issue stage coupled to the decoding component to receive the microinstruction sequence perform optimization processing by reordering the microinstruction sequence into an optimized microinstruction sequence comprising a plurality of dependent code groups;

a microprocessor pipeline coupled to the allocation and issue stage to receive and execute the optimized microinstruction sequence; and

a sequence cache coupled to the allocation and issue stage to receive and store a copy of the optimized microinstruction sequence for subsequent use upon a subsequent hit on the optimized microinstruction sequence.

9. The microprocessor of claim 8, wherein a copy of the decoded microinstructions are stored in a microinstruction cache.

10. The microprocessor of claim 8, wherein the optimization processing is performed using an allocation and issue stage of the microprocessor.

11. The microprocessor of claim 10, wherein the allocation and issue stage further comprises an instruction scheduling and optimizer component that reorders the microinstruction sequence into the optimized micro instruction sequence.

12. The microprocessor of claim 8, wherein the optimization processing further comprises dynamically unrolling microinstruction sequences.

13. The microprocessor of claim 8, wherein the optimization processing is implemented through a plurality of iterations.

14. The microprocessor of claim 8, wherein the optimization processing is implemented through a register renaming process to enable the reordering.

15. In a microprocessor, a method for accelerating code optimization, comprising:

accessing an input microinstruction sequence by using a software-based optimizer instantiated in memory;

using SIMD instructions to populate a dependency matrix with dependency information extracted from the input microinstruction sequence;

scanning a plurality of rows of the dependency matrix to perform optimization processing by reordering the microinstruction sequence into an optimized microinstruction sequence comprising a plurality of dependent code groups;

outputting the optimized microinstruction sequence to a microprocessor pipeline for execution; and

storing a copy of the optimized microinstruction sequence into a sequence cache for subsequent use upon a subsequent hit optimized microinstruction sequence.

16. The method of claim 15, wherein optimization processing further includes scanning the plurality of rows of the dependency matrix to identify matching instructions.

17. The method of claim 16, wherein optimization processing further includes analyzing the matching instructions to determine whether the matching instructions comprise a blocking dependency, and wherein renaming is performed to remove the blocking dependency.

18. The method of claim 17, wherein instructions corresponding to first matches of each row of the dependency matrix are moved into a corresponding dependency group.

19. The method of claim 15, wherein copies of the optimized microinstruction sequences are stored in a memory hierarchy of the microprocessor.

20. The method of claim 19, wherein the memory hierarchy comprises an L1 cache and an L2 cache.

21. The method of claim 20, wherein the memory hierarchy for further comprises a system memory.

22. A microprocessor, comprising:

an instruction fetch component for fetching an incoming microinstruction sequence;

a decoding component coupled to the instruction fetch component to receive the fetched macro instruction sequence and decode into a microinstruction sequence;

an allocation and issue stage coupled to the decoding component to receive the microinstruction sequence perform optimization processing by reordering the microinstruction sequence into an optimized microinstruction sequence comprising a plurality of dependent code groups;

a microprocessor pipeline coupled to the allocation and issue stage to receive and execute the optimized microinstruction sequence;

a sequence cache coupled to the allocation and issue stage to receive and store a copy of the optimized microinstruction sequence for subsequent use upon a subsequent hit on the optimized microinstruction sequence; and

a hardware component for moving instructions in the incoming microinstruction sequence.

23. The microprocessor of claim 22, wherein at least one register is renamed and at least one instruction is moved ahead of the branch without inserting compensation code.

24. The microprocessor of claim 23, wherein the hardware component keeps track of whether a branch biased decision is true, and wherein in case of a wrongly predicted branch, the hardware component automatically rolls back state in order to execute a correct instruction sequence.

25. The microprocessor of claim 24, wherein the hardware component jumps to original code in memory to execute the correct instruction sequence in case of a wrongly predicted branch.

26. The microprocessor of claim 25 wherein the hardware component causes a flushing of a miss predicted instruction sequence in the case of a wrongly predicted branch.

* * * * *