(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
19 February 2004 (19.02.2004)

PCT

(10) International Publication Number
WO 2004/015562 A2

(51) International Patent Classification[7]: G06F 9/38

(21) International Application Number: PCT/US2003/024121

(22) International Filing Date: 31 July 2003 (31.07.2003)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
10/215,095 8 August 2002 (08.08.2002) US

(71) Applicant: **HEWLETT-PACKARD DEVELOPMENT COMPANY, L.P.** [US/US]; 20555 S.H. 249, Houston, TX 77070 (US).

(72) Inventor: **SCHLANSKER, Michael S,**; 409 La Prenda, Los Altos, CA 94024 (US).

(74) Agent: **LEE, Denise A,**; Hewlett-Packard Company, Intellectual Property Administration, P O Box 272400, Mail Stop 35, Fort Collins, CO 80527-2400 (US).

(81) Designated States *(national)*: AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VC, VN, YU, ZA, ZM, ZW.

(84) Designated States *(regional)*: ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, RO, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:
— *without international search report and to be republished upon receipt of that report*

[Continued on next page]

(54) Title: SYSTEM AND METHOD USING DIFFERENTIAL BRANCH LATENCY PROCESSING ELEMENTS

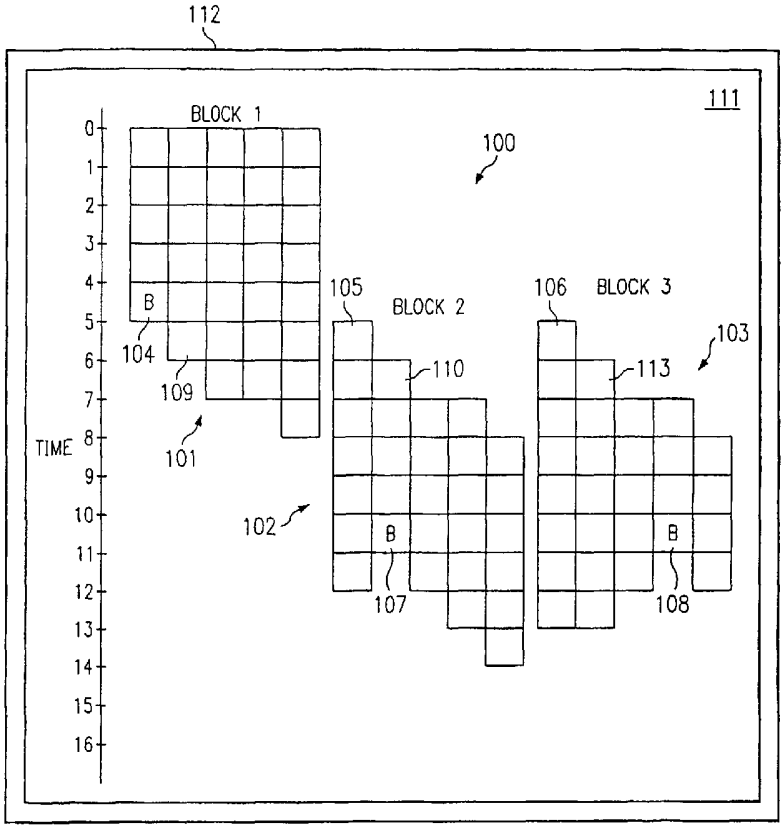

(57) Abstract: The invention is a system and method for executing a program that comprises a plurality of basic blocks on a computer system that comprises a plurity of processing elements (202). The invention generates a branch instruction by one processing element (202) of the plurality of processing elements (202), sends the branch instruction to the plurality of processing elements (202). The invention then independently branches to a target of the branch instruction by each of the processing elements (202) of the plurality of processing elements (202) when each processing element receives the sent branch instruction. At least one processing element (202) of the plurality of processing elements (202) receives the branch instruction at a time later than another processing element (202) of the plurality of processing elements (202).

WO 2004/015562 A2

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

SYSTEM AND METHOD USING DIFFERENTIAL BRANCH LATENCY PROCESSING ELEMENTS

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] The present application is related to copending and commonly assigned United States patent application serial number [Attorney Docket No. 100110353-1] entitled "Branch Reconfigurable Systems and Methods," filed concurrently herewith, the disclosure of which is hereby incorporated by reference.

FIELD OF THE INVENTION

[0002] This invention relates in general to computers, and in specific to a system and method that permits variable latency in branching operations.

BACKGROUND OF THE INVENTION

[0003] Typical general purpose computer systems comprise one of many different architectures. Architecture, as used herein, refers to the instruction set and resources available to a programmer for a particular computer system. Thus, architecture includes instruction formats, instruction semantics, operation definitions, registers, memory addressing modes, address space characteristics, etc. An implementation is a hardware design or system that realizes the operations specified by the architecture. The implementation determines the characteristics of a microprocessor that are most often measured, e.g. price, performance, power consumption, heat dissipation, pin number, operating frequency, etc. Thus, a range of implementations of a particular architecture can be built, but the architecture influences the quality and cost-effectiveness of those implementations. The influence is exerted largely in the trade-offs that must be made to accommodate the complexity associated with the instruction set.

[0004] Most architectures try to increase efficiency in their respective implementations by exploiting some form of parallelism. For example, in single instruction multiple data stream (SIMD) architecture implementations, the various processing elements (PEs) can all perform the same operation at the same time, each with its own local (different) data.

[0005] One common architecture is the very long instruction word (VLIW) architecture. Although very similar to SIMD systems, in VLIW system, each PE can perform a different operation independent of the other PEs. However, the grouping of the sets of operations that PEs can execute together is static. In other words, the choice of which operations that can simultaneously execute together is made at compile time. Moreover, their execution is synchronous. This means that each of the PEs is processing the instructions in a lock-step manner. Note that VLIW PEs are sometimes referred to as function units (FUs), because some PEs within a VLIW system may support only certain types of operations.

[0006] VLIW processors are wide-issue processors that use static scheduling to orchestrate the execution of a number of parallel processor elements. VLIW processors have constant branch latency. When a branch is executed on one of a multiplicity of processing elements, the effect of the branch occurs simultaneously on all processor elements. That is, if a branch issued on the t-th cycle and the branch latency is q cycles, all function units begin execution of code at the branch target at cycle t+q.

[0007] VLIW processors use a program counter to index into an instruction memory to select an instruction that is used to simultaneously control all processing elements. Since the instruction is taken from the instruction memory in a single atomic action, program text that is taken from the instruction memory after a branch is available to all function units on the same program cycle.

[0008] Since the processing elements are physically separate from each other, a branch that occurs in an originating processing element may take different amounts of time to reach each of the processing elements. VLIW scheduling requires that the branch would have to take effect at all processing elements at the same time. Thus, the branch command at the closest processing element would have to be delayed for a time equal to the branch delay to the farthest processing element.

[0009] FIGURE 4 depicts a conventional scheduling model that has uniform latency. Basic blocks of program code are labeled block 1 401, block 2 402, and block 3 403. Increasing time is represented by moving downward one row for each clock cycle. This VLIW system is a 5-way VLIW processor and has five processing elements operating on the basic blocks, as represented by the columns. This system has a branch latency of 3 cycles, thus it

takes 2 additional cycles after a branch is encountered for all of the processing elements to receive the branch. As shown in FIGURE 4, the first processing element forms a conditional branch B 404 while processing the 6th cycle of basic block 1 401. This conditional branch may lead to processing of block 2 402 or block 3 403, depending upon whether the condition is satisfied or not. For example, block 2 402 may be processed if the condition is not satisfied, i.e. falls through, and block 3 403 may be processed if the condition is met, i.e. branch taken. Note that the fall-though block is normally contiguous in memory with the prior block from which the branch was issued.

[0010] In any event, the first processing element cannot immediately begin execution on either block 2 or block 3, but rather must wait for 2 cycles until all of the other processing elements are ready to move to the next block. This two cycle branch delay causes a gap between the location of the branch in the code and the actual end of the basic block. This gap is called the branch shadow. Operations within the two cycle branch shadow execute unconditionally as if the branch has not yet been executed. Useful operations which should execute irrespective of the branch condition or no-ops 406 may be executed within the branch shadow, however, for any operation that should not execute when the branch is taken, the operation must appear below this two cycle window.

[0011] After waiting, all of the processing elements then move to either block 2 or block 3. For example, the first processing element starts execution at either location 405 or 407, of blocks 2 or 3, respectively. During processing of block 2 or block 3, another branch would be encountered, e.g. branch 408 or 409, which would change the flow of the program to other basic blocks (not shown). Again, because of the three cycle branch latency, the branch originating processing element would wait for two additional cycles before processing the subsequent basic block so that the other processing elements would have received the branch.

[0012] A problem with this arrangement is the cycles lost to waiting for the branch latency. For programs with a great deal of branches, these lost cycles can greatly reduce the efficiency of the system.

BRIEF SUMMARY OF THE INVENTION

[0013] An embodiment of the invention is a system and method for executing a program that comprises a plurality of basic blocks on a computer system that comprises a

plurality of processing elements. This invention generates a branch instruction by one processing element of the plurality of processing elements and sends the branch instruction to the plurality of processing elements. This invention embodiment then independently branches to a target of the branch instruction by each of the processing elements of the plurality of processing elements when each processing element receives the sent branch instruction. At least one processing element of the plurality of processing elements receives the branch instruction at a time later than another processing element of the plurality of processing elements.

BRIEF DESCRIPTION OF THE DRAWINGS

[0014] FIGURE 1 is block diagram depicting an example of a program schedule formed by a compiler, according to an embodiment of the invention.

[0015] FIGURE 2 is a block diagram of an example of a computer system operating with the schedule of FIGURE 1, according to an embodiment of the invention.

[0016] FIGURE 3 depicts a flow chart showing aspects of the compilation and execution of a program, according to an embodiment of the invention.

[0017] FIGURE 4 is a block diagram of an example of a prior art program schedule.

DETAILED DESCRIPTION

[0018] The invention uses a computer architecture that supports the VLIW mode of operation, but preferably relaxes the timing constraint in processing branch executions. Thus, each processing element may preferably begin processing the branch to the next basic block as soon as it is received. Thus, latency time for the time of flight for the branch instructions does not have to be padded so that all processing elements begin the next basic block simultaneously. During compilation, the compiler may model the program and the hardware in such a way that the differential flight time with regards to the processing of the basic blocks is accounted for in the schedule. Thus, the compiler can map the scheduling of the program with respect to the flight time of branch instructions and the destinations of the branches. Such flight time can preferably be tabulated by the compiler, and represented as a table of vectors that defines time from an originating processing element to the destination processing elements, which preferably includes the originating processing element. The hardware would then be able to allow

processing elements to branch as soon as the branch is received, rather than padding branch flight time into the processing elements. Thus, the invention realizes higher performance than a conventional VLIW processor.

[0019] The inventive differential branch latency VLIW preferably allows for different processing elements to respond to a single branch command at different points in time. A closer processing element may respond more quickly than a farther processing element. The invention preferably allows compilers to accommodate non-uniform scheduling models. This allows processing to begin within a new basic block at lower latency within certain processing elements than would otherwise be possible if all latencies were padded to a maximal latency.

[0020] FIGURE 1 depicts a block diagram of an example of a schedule 100 of the processing of basic blocks according to one aspect of the invention. The scheduling is performed by compiler 111 operating on computer 112. In referring to FIGURE 3, during compilation 31, the schedule for the program is developed 33. Basic blocks of program code are labeled block 1 101 , block 2 102, and block 3 103. Increasing time is represented by moving downward one row for each clock cycle. The system is preferably a VLIW system, and as shown is a 5-way VLIW processor and has five processing elements operating on the basic blocks, as represented by the columns. This system has a maximum branch latency of 4 cycles, thus it takes 3 additional cycles after a branch is encountered for all of the processing elements to receive the branch.

[0021] Referring to FIGURE 3, which depicts an example 30 of the operation of the invention during compilation and execution of a program, the invention begins execution of the program in step 34, which would include the basic blocks of FIGURE 1. As shown in FIGURE 1, the first processing element forms a conditional branch B 104 while processing the 5^{th} cycle of basic block 1 101. Referring to FIGURE 3, this occurs in step 35. This conditional branch may lead to processing of block 2 102 or block 3 103, depending upon whether the condition is satisfied or not. For example, block 2 102 may be processed if the condition is not satisfied, i.e. falls through, and block 3 103 may be processed if the condition is met, i.e. branch taken. Note that the fall-though block is normally contiguous in memory with the prior block from which the branch was issued.

[0022] The branching processor would send the branch to the other processors, as indicated by step 36 of FIGURE 3. The invention allows the first processing element to immediately move to either block 2 or block 3, and thus, does not have to wait 3 cycles, the time needed to branch from the first processing element to the fifth processing element, or until all of the other processing elements have received the branch. Other processing elements will move to the appropriate block upon receipt of the branch. For example, the second processing element moves from block 109 to either block 2 110 or block 3 113 one cycle after the first processing element. Thus, the time at which block 2 is reached is not identical to, and is independent of, the time that block 3 is reached. This independent branching by the processors is shown as step 37 of FIGURE 3.

[0023] Within 4 cycles, all of the processing elements have moved to either block 2 or block 3 and commenced executing the code of the branched to block, as indicated by step 38 of FIGURE 3. For example, the first processing element starts execution at either location 105 or 106, of blocks 2 or 3, respectively. During processing of block 2 or block 3, another branch would be encountered, e.g. branch 107 or 108, which would change the flow of the program to other basic blocks (not shown). The processing elements would immediately move to the other blocks after receiving the branch instruction. Any of these other basic blocks may contain a branch that causes the program to generate a branch 35 again repeating the process until the program terminates. The processing elements would continue with the execution of the program to the logical end of the program, as shown by step 39 of FIGURE 3. In some systems, the number of processing elements and their branch latencies may change during program execution. In such cases, on-line recompilation is used as the program may return to the compilation step 31, regenerate branch latency tables, and resume execution of a program that is recompiled with a new latency configuration.

[0024] The branch latency preferably is represented by a variable, e.g. a branch latency vector for each of the processing elements that might originate a branch. This vector would describe the time for a branch to move from an originating processing element to each of the destination processing elements. In the example shown in FIGURE 1, the branch latency vector (BLV(j)) for the first processing element would be BLV(1)=(1,2,3,3,4). For the second processing element, the vector would be BLV(2)=(2,1,2,3,4), and for the fourth processing element the vector would be BLV(4)=(3,3,2,1,2). Each entry in the vector indicates the latency

of the corresponding column with respect to issuing a branch in the jth column. Note that it is assumed that every branch takes a shortest path through the branch transport network and the traversal of each branch transport node takes one cycle. In general, each processing element j, that executes a branch, will have a distinct branch latency vector BLV(j). The vector defines the distinct number of branch delay slots or cycles for each of the processing elements for a branch from the originating processing element. Thus, the vector defines the start times for the processing elements to begin processing the branched-to basic block. The compiler preferably forms the vectors, and preferably stores the vectors in a branch latency table. The compiler preferably forms the vectors prior to (or contemporaneous to) scheduling. Referring to FIGURE 3, this is shown in step 32.

[0025] The vector for the first processing element is used to structure the code for the branch in block 1, while the vectors for the second and fourth processing elements are used to structure code for the branches in blocks 2 and 3, respectively. Note that the vector affects the ending of the current basic block and the beginning of the next basic block or blocks. Thus, the branch 104 affects the ending of block 1 and the beginnings of blocks 2 and 3, while branch 107 affects the ending of block 2 (and the beginning of a block that is not shown). Similarly, branch 108 affects the ending of block 3 (and the beginning of a block that is not shown). Note that the beginnings of the branched-to blocks, e.g. blocks 2 and 3, match the ending of the branching block, e.g. block 1. Note that, in this example machine, the originating processing element can respond to a branch on the next cycle, so the originating processing element has a latency of 1.

[0026] A compiler, using the vectors, can schedule the execution of the program with respect to the processing elements. Referring to FIGURE 3, this is shown in step 33. For example, the compiler can schedule work for basic block 1 to be performed by the second processing element for up to one cycle after the branch 104. However, after that one cycle, the compiler should then schedule only work for the second or third basic blocks. So in this manner, the compiler may create a static schedule which carefully acknowledges when the branches occur, and where processing is to be assigned within the basic blocks to accommodate the branches.

[0027] FIGURE 2 depicts an example of an arrangement of a branch transport network 201 for a plurality of processing elements 202-1 to 202-N that can be used with the schedule of FIGURE 1. Note that FIGURE 1 uses five processing elements, while FIGURE 2

depicts more processing elements. The additional processing elements may be used to handle other programs. Additional information on configuring the branch transport cells and the processor elements is described in related application entitled "Branch Reconfigurable Systems and Methods," [Attorney Docket No. 100110353-1] and hereby incorporated herein by reference. Note that the hardware in embodiments of the present invention preferably do not include the latency buffers of this related application.

[0028] Each of the processing elements 202 includes an instruction memory for holding instructions to be processed by its function unit. When a branch command is executed by one of the processing elements, the name of the branch-to basic block is transmitted to the desired processing elements. This named program location is translated (e.g. by content addressable RAM or table lookup) to potentially distinct actual program locations within each of the separate instruction memories for each of the branch units. After the new branch target is reached, each of the processing elements begins independently sequencing through the code from the branch-to basic block.

CLAIMS

What is claimed is:

1. A method for executing a program that comprises a plurality of basic blocks on a computer system that comprises a plurality of processing elements (202):

generating a branch instruction by one processing element (202) of the plurality of processing elements (202);

sending the branch instruction to the plurality of processing elements (202); and

independently branching to a target of the branch instruction by each of the processing elements (202) of the plurality of processing elements (202) when each processing element (202) receives the sent branch instruction;

wherein at least one processing element (202) of the plurality of processing elements (202) receives the branch instruction at a time later than another processing element (202) of the plurality of processing elements (202).

2. The method of claim 1, further comprising:

after branching, executing a basic block located at the target by each processing element (202) of the plurality of processing elements (202).

3. The method of claim 1, wherein the sending comprises:

transporting the branch instruction by a network.

4. The method of claim 1, further comprising:

forming a branch latency table for the system, wherein the table comprises a plurality of branch vectors that describe a latency for sending a branch from the one processing element (202) to each processing element of the plurality of processing elements (202).

5. The method of claim 1, wherein:

the branch is a conditional branch that has at least two targets, wherein selection of each target is dependent upon satisfaction of a condition.

6. A system for executing a program that comprises a plurality of basic blocks, the system comprising:

a plurality of processing elements (202), wherein at least one processing element (202) of the plurality of processing elements generates a branch instruction during processing of a basic block; and

a branch transport network (201) that delivers the branch instruction to at least one other processing element (202) of the plurality of processing elements;

wherein the at least one processing element (202) and the at least one other processing element (202) branch to a target of the branch instruction independently of each other; and

wherein the at least one other processing element (202) receives the branch instruction at a time different from the one processing element (202).

7. The system of claim 6, wherein:

the computer system is a VLIW system.

8. The system of claim 6, further comprising:

a branch latency table having at least one branch vector that describe a latency for sending a branch from the at least one processing element (202) to the at least one other processing element (202).

9. The system of claim 6, wherein:

the program is statically scheduled by a compiler.

10. The system of claim 6, wherein:

the branch is a conditional branch that has at least two targets, wherein selection of each target is dependent upon satisfaction of a condition.

FIG. 1

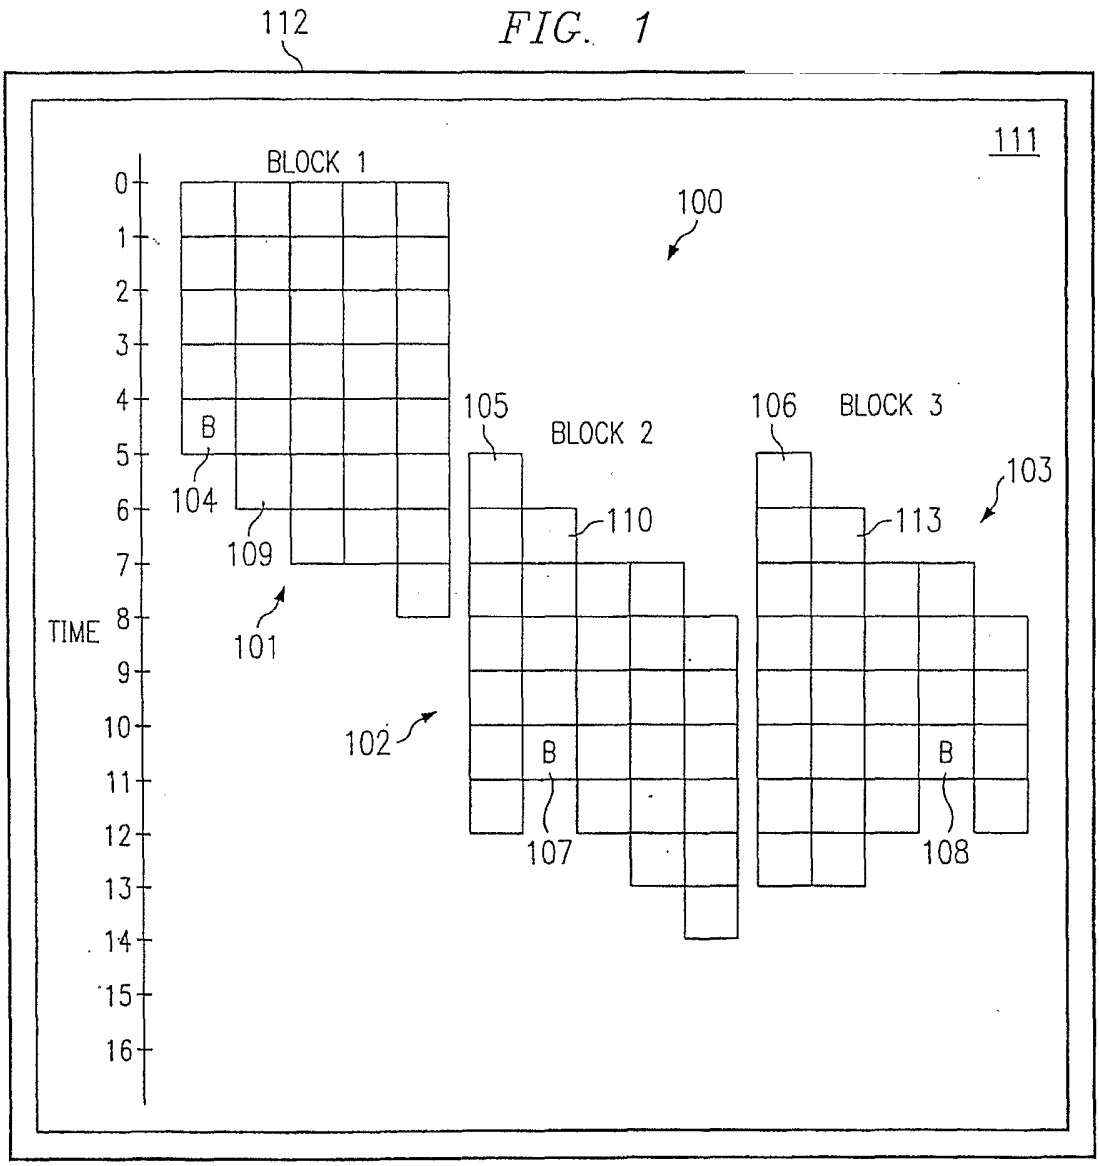
112
111
BLOCK 1
BLOCK 2
BLOCK 3
100
105
106
103
110
113
B
104
109
101
102
107
108
TIME
0
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16

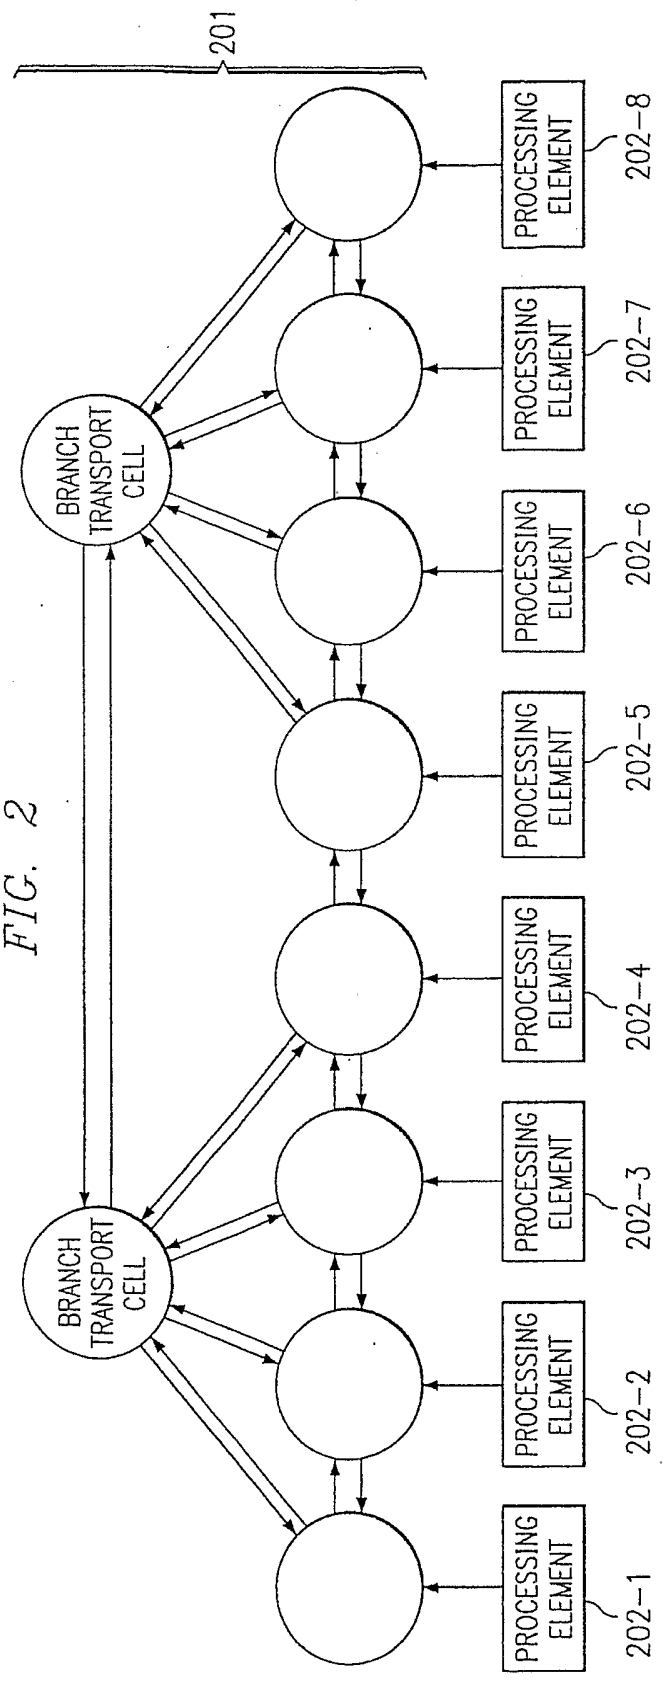
201
BRANCH TRANSPORT CELL
BRANCH TRANSPORT CELL
PROCESSING ELEMENT
202-1
PROCESSING ELEMENT
202-2
PROCESSING ELEMENT
202-3
PROCESSING ELEMENT
202-4
PROCESSING ELEMENT
202-5
PROCESSING ELEMENT
202-6
PROCESSING ELEMENT
202-7
PROCESSING ELEMENT
202-8

FIG. 2

FIG. 3

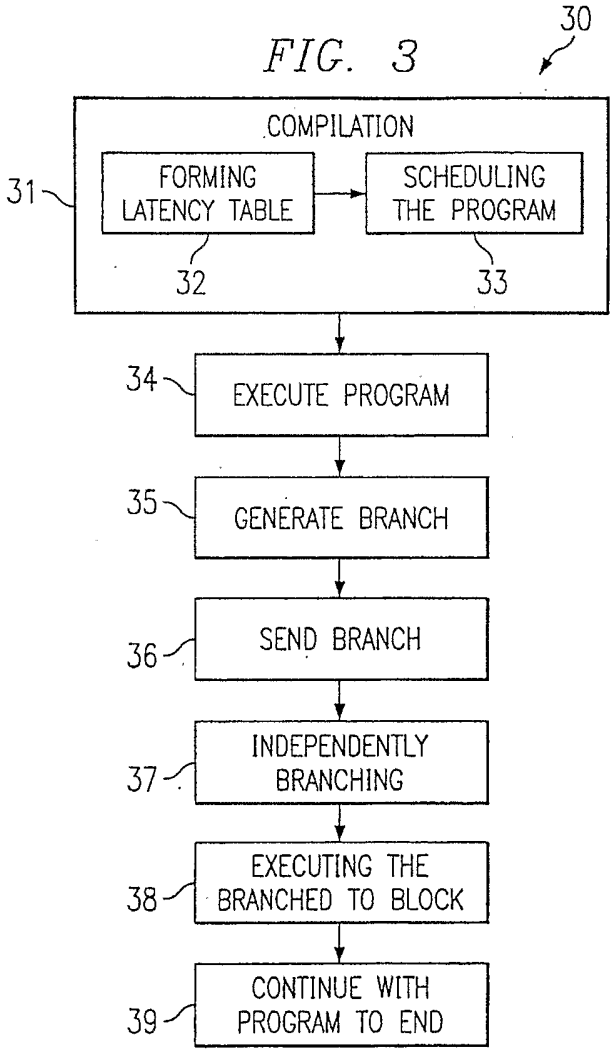
30
COMPILATION
31
FORMING LATENCY TABLE
SCHEDULING THE PROGRAM
32
33
34
EXECUTE PROGRAM
35
GENERATE BRANCH
36
SEND BRANCH
37
INDEPENDENTLY BRANCHING
38
EXECUTING THE BRANCHED TO BLOCK
39
CONTINUE WITH PROGRAM TO END

FIG. 4
(PRIOR ART)