

(51) International Patent Classification:
G06F 9/06 (2006.01) **G06F 9/38** (2006.01)(21) International Application Number:
PCT/US2011/055917(22) International Filing Date:
12 October 2011 (12.10.2011)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/392,391 12 October 2010 (12.10.2010) US(71) Applicant (for all designated States except US): **SOFT MACHINES, INC.** [US/US]; 3211 Scott Boulevard, Suite 202, Santa Clara, CA 95054 (US).

(72) Inventor; and

(75) Inventor/Applicant (for US only): **ABDALLAH, Mo-hammad** [US/US]; 3868 Suncrest Avenue, San Jose, CA 95132 (US).(74) Agent: **BARNES, Glenn D.**; Murabito Hao & Barnes LLP, Two North Market Street, Third Floor, San Jose, CA 95113 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

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

Published:

— without international search report and to be republished upon receipt of that report (Rule 48.2(g))

(54) Title: AN INSTRUCTION SEQUENCE BUFFER TO ENHANCE BRANCH PREDICTION EFFICIENCY




FIGURE 3

(57) Abstract: A method for outputting alternative instruction sequences. The method includes tracking repetitive hits to determine a set of frequently hit instruction sequences for a microprocessor. A frequently miss-predicted branch instruction is identified, wherein the predicted outcome of the branch instruction is frequently wrong. An alternative instruction sequence for the branch instruction target is stored into a buffer. On a subsequent hit to the branch instruction where the predicted outcome of the branch instruction was wrong, the alternative instruction sequence is output from the buffer.

AN INSTRUCTION SEQUENCE BUFFER TO ENHANCE
BRANCH PREDICTION EFFICIENCY

FIELD OF THE INVENTION

[001] 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

[002] Improving computer architecture performance is a difficult task. Improvements have been sought through frequency scaling, Single Instruction Multiple Data (SIMD), Very Long Instruction Word (VLIW), multi-threading and multiple processor techniques. These approaches mainly target improvements in the throughput of program execution. Many of the techniques require software to explicitly unveil parallelism. In contrast, frequency scaling improves both throughput and latency without requiring software explicit annotation of parallelism. Recently, frequency scaling hit a power wall so improvements through frequency scaling are difficult. Thus, it is difficult to increase throughput unless massive explicit software parallelization is expressed.

[003] With respect to single threaded program execution, program execution is controlled by branching instructions that dictate the program control flow. Program instruction sequences are dynamic when the branching instructions are conditional or the branch target is indirect. In such cases, it is essential for the fetch logic of the processor to find out for conditional branches if the branch is taken or not taken. This enables the fetch logic to bring in the sequence of instructions that either follow the target of the branch or those that follows the branch instruction itself. There exists a problem, however, in that at the fetch stage, the outcome of the condition of the branch is not known before the branch itself executes.

[004] In an attempt to overcome this problem, prior art designs have implemented branch prediction logic to predict the outcome of a branch. At the fetch stage of the microprocessor, the predicted outcome enables the fetch logic to anticipate where to bring the next sequence of instructions from. Problems still exists, however, since this processing needs to be sequential in nature. The current branch needs to be processed first in order to know where to bring the next instruction sequence. Accordingly the sequential nature of processing branches in the fetch stage imposes a performance bottleneck on the single threaded execution speed of a microprocessor. Penalties for an incorrect branch prediction typically involve flushing the whole pipeline of a microprocessor, accessing caches and reloading with a new instruction sequence. These penalties greatly reduce the incentives for predicting more than one branch at a time.

SUMMARY OF THE INVENTION

[005] Embodiments of the present invention implement an algorithm (e.g., a method and an apparatus) that increases the efficiency of branch production processing of instruction sequences.

[006] In one embodiment, the present invention is implemented as a method for outputting alternative instruction sequences. The method includes tracking repetitive hits to determine a set of frequently hit instruction sequences for a microprocessor. Where in a frequently miss-predicted branch instruction is identified, where in the predicted outcome of the branch instruction is frequently wrong. An alternative instruction sequence for the branch instruction target is stored into a buffer. On a subsequent hit to the branch instruction where the predicted outcome of the branch instruction was wrong, the alternative instruction sequence is output from the buffer.

[007] 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

[008] 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.

[009] Figure 1 shows an exemplary sequence of instructions operated on by one embodiment of the present invention.

[010] Figure 2 shows the sequence instructions with the respective code segments for each branch illustrated in accordance with one embodiment of the present invention.

[011] Figure 3 shows a flow diagram of an apparatus used to output alternative instruction sequences for branches that are frequently hit and are frequently miss-predicted in accordance with one embodiment of the present invention.

[012] Figure 4 shows an overview flowchart of the steps of a process for outputting alternative instruction sequences in accordance with one embodiment of the present invention.

[013] Figure 5 shows a diagram of an instruction sequence buffer in accordance with one embodiment of the present invention.

[014] Figure 6 shows a diagram of an instruction sequence buffer that is used to store instruction sequences for reliably predictable branches that are frequently hit in accordance with one embodiment of the present invention.

[015] Figure 7 shows an overview flowchart of the steps of a process for outputting reliably predictable instruction sequences in accordance with one embodiment of the present invention.

[016] Figure 8 shows a diagram of an exemplary microprocessor pipeline in accordance with one embodiment of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

[017] 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.

[018] 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.

[019] 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.

[020] 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.

[021] 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.

[022] In one embodiment, the present invention implements an algorithm (e.g., a method and an apparatus) for outputting alternative instruction sequences for branches that are frequently hit and are frequently miss-predicted. The method includes tracking repetitive hits to a branch instruction to determine a set of frequently hit instruction sequences for a microprocessor. Then frequently miss-predicted branch instructions are identified, wherein the predicted outcome of the branch instruction is frequently wrong. An alternative instruction sequence for the branch instruction is stored into a buffer (e.g., an instruction sequence buffer). On a subsequent hit to the branch instruction where the predicted outcome of the branch instruction was wrong, the alternative instruction sequence is output from the buffer. The alternative instruction sequence thus saves the microprocessor pipeline from being flushed in its entirety. The alternative instruction sequence is provided directly from the buffer, as opposed to, for example, flushing the whole pipeline, accessing caches and assembling a new instruction sequence. Figure 1 shows an exemplary sequence of instructions operated on by embodiments of the present

invention. Subsequently, Figure 2 shows a flow diagram of alternative instruction sequences produced by multiple levels of branches, and Figure 3 shows an overview flowchart of the steps of a process for outputting alternative instruction sequences in accordance with one embodiment of the present invention.

[023] In an alternative embodiment, dual use of the storage resources of the instruction sequence buffer is implemented. Instead of storing alternative instruction sequences for frequently hit in frequently miss predicted branches, the storage resources of the instruction sequence buffer are used to store instruction sequences for frequently hit and reliably predicted branches. Thus, instead of storing alternative instruction sequences for the taken case and the not taken case, the storage resources of the buffer 600 are used to store the instruction sequences of a frequently hit and reliably predictable branch and a number of the subsequent following branches. This alternative embodiment is shown and described in Figure 6 below. The two embodiments can both exist together and share the same storage buffer but in a different manner.

[024] Figure 1 shows an exemplary sequence of instructions operated on by one embodiment of the present invention. As depicted in Figure 1, the instruction sequence 100 comprises 16 instructions, proceeding from the top of Figure 1 to the bottom. As can be seen in Figure 1, the sequence 100 includes four branch instructions 101-104.

[025] One objective of embodiments of the present invention is to output alternative instruction sequences for branches that are frequently hit and are frequently miss-predicted. The alternative instruction sequences are output as a means of greatly reducing the latency penalty of assembling an alternative instruction sequence from scratch. In accordance with different embodiments, these instructions can comprise native instructions (e.g., native instructions of the microprocessor architecture, such as x86 instructions, MIPS instructions, or the like). Alternatively, these instructions can comprise microcode. As described earlier, the more branches a sequence of instructions include, the more combinations and possible resulting sequences occur and need to be dealt with. This characteristic is illustrated in Figure 2 below.

[026] Figure 2 shows the sequence instructions 100 with the respective code segments for each branch illustrated in accordance with one embodiment of the present invention. As described above, the more branches that are presented in a sequence of instructions, the more combinations and possibilities of sequences of instructions that need to be disambiguated.

[027] This is illustrated in Figure 2, which shows a first resulting sequence “1” that occurs if branch c1 is taken. As referred to herein, a branch is taken if the program execution flow moves to the target of the branch. This is indicated by the two digits within parenthesis at the end of each of the branch instructions. For example, branch c1 has a target of 11 and results in skipping the next 6 instructions. Similarly, branch c2 has a target of 10 and results in skipping the next 2 instructions, and so on.

[028] Thus, a second resulting sequence “2” is shown, and occurs if branch c2 is taken. A third resulting sequence “3” is shown as occurring if branch c3 is taken. Similarly, the fourth resulting sequence “4” is shown as occurring if branch c4 is taken.

[029] Embodiments of the present invention output alternative instruction sequences for branches that are frequently hit and are frequently miss-predicted. As shown in Figure 2, a different instruction sequence occurs when any of the branches along the predicted instruction sequence is miss-predicted. Embodiments of the present invention advantageously store a number of these alternative instruction sequences in a buffer that is located very close to the decoder hardware. The stored alternative instruction sequences are output as a means of greatly reducing the latency penalty of assembling an alternative instruction sequence from scratch. This algorithm is further diagrammed below in Figure 3.

[030] Figure 3 shows a flow diagram of an apparatus 300 used to output alternative instruction sequences for branches that are frequently hit and are frequently miss-predicted in accordance with one embodiment of the present invention. As depicted in Figure 3, apparatus 300 includes a sequence predictor 301, a branch prediction table 302, and a sequence of stability counter 303.

[031] In the Figure 3 embodiment, the apparatus 300 functions by tracking repetitive hits to determine a set of frequently hit branches and their corresponding instruction sequences. These branches are illustrated as B0 through B8 in Figure 3. As described above, a predicted instruction sequence is assembled based upon the branch predictions for the branches. Lines are shown connecting the first branch B0 to the following branches B1 and B5, and from B1 and B5 to their respective following branches B2, B4 and B6, and so on.

[032] The branch prediction table 302 is used to keep track of the branch predictions in order to determine which instruction sequences are often miss predicted and which instruction sequences are very rarely miss predicted.

[033] In the Figure 3 embodiment, branches are analyzed three levels deep past the initial branch B0. Thus, for example, an alternative instruction sequence can be assembled from B0 to B1, to B2 and on to B3. Depending upon the size of the buffer, a greater or lesser number of levels of following branches can be analyzed and stored.

[034] The sequence predictor 301 functions by predicting the outcome of the branches to assemble a predicted instruction sequences. Accordingly, the sequence predictor can monitor the execution of these branches and identify frequently miss predicted branch instructions. A different number of mechanisms can be used to identify the frequently miss predicted branch instruction. In one embodiment, a tag structure such as a sequence stability counter 303 is used to accumulate repetitive hits to the same instruction sequences (e.g., frequently fetched sequences). Once a threshold has been exceeded, a given branch instruction can be identified as and treated as a frequently miss predicted branch instruction.

[035] Figure 4 shows an overview flowchart of the steps of a process 400 for outputting alternative instruction sequences in accordance with one embodiment of the present invention. Process 400 shows exemplary operating steps of, for example, an instruction fetch module of a microprocessor.

[036] Process 400 begins in step 401, where accesses are tracked in order to determine a set of frequently hit instruction sequences. As described above, the sequence predictor assembles predicted instruction sequences. It can determine which instruction sequences are frequently fetched.

[037] In step 402, out of that set, the sequence predictor identifies a frequently miss-predicted branch instruction. As described above, the predicted outcome of this branch instruction is frequently wrong.

[038] In step 403, an alternative instruction sequence for the branch instruction is stored into a buffer. As described above, this inability to correctly and reliably predict the outcome of this branch could lead to frequent pipeline flushes. However, in accordance with embodiments of the present invention, alternative instruction sequences are stored within an instruction sequence buffer.

[039] In step 404, on a subsequent hit to the branch instruction where the predicted outcome was wrong, the alternative instruction sequence is output from the buffer. The alternative instruction sequence thus saves the entire microprocessor pipeline from being flushed. The alternative instruction sequence is provided directly from the buffer, as opposed to, for example, flushing the whole pipeline, accessing caches and assembling a new instruction sequence.

[040] It should be noted that in one embodiment, the branch prediction table can be used to prequalify the instruction sequences that enter into the tag structure using a saturating counter in the branch prediction table to count number of repetitive accesses to a certain branch. Once that reaches saturation then the instruction sequence address that starts at the outcome of that branch is entered into the tracking tag structure. This structure has a smaller number of entries and has larger saturating counters. Once the counter reaches a threshold of a large count that justifies storing the instruction sequence, then that instruction sequence is entered into the buffer.

[041] It should be noted that in one embodiment, a different threshold might be used if the sequences that are hit are based on well predicted branches or frequently miss

predicted branches. Thereafter, the instruction sequence is accessed from this buffer every time the fetch hardware jumps to this particular instruction sequence and is indexed by the address of the branch that leads to it.

[042] Figure 5 shows a diagram of an instruction sequence buffer 500 in accordance with one embodiment of the present invention. As depicted in Figure 5, the buffer 500 includes three portions 501-503. The buffer 500 and the portions 501-503 show an example of embodiments of the present invention storing alternative instruction sequences for each of the possible following branches that flow from branch B0. For each of the branches B1 through B8, the possible resulting instruction sequences from each branch being taken or not taken are stored. For example, instruction sequences for branch B0 being taken (e.g., leading to B1) or not taken (e.g., leading to B5) are stored into the buffer 500. Similarly, instructions for branch B1 being taken (e.g., leading to B2) are not taken (e.g., leading to B4) are stored into the buffer 500, and so on for each of the following branches.

[043] In this manner, the buffer 500 includes all the possible instruction sequences that flow from branch B0. This attribute allows a very quick recovery from a miss prediction of the branch B0. For example, on a subsequent hit to the branch instruction B0, where the predicted outcome of the branch instruction was wrong, an alternative instruction sequence can be rapidly output from the buffer 500. This avoids the necessity of flushing the whole pipeline, accessing caches and assembling a new alternative instruction sequence from scratch (e.g., re-fetching the instructions, etc.).

[044] The Figure 500 embodiment shows how the portions 501-503 include instruction sequences for both the taken case and the not taken case for each of the branches B0 through B8. For example, portion 501 shows instructions for the taken case stored on a first way of the portion on the left hand side. This is illustrated by the "T" at the top of the portion. Instructions for the not taken case are stored on the right hand side, as illustrated by the "NT" at the top of the portion. The taken and not taken cases represent two ways into which the buffer portion, or cache, can be indexed. This is illustrated as

way 1 “W1” and way 2 “W2” at the top of the portion. These attributes are similarly illustrated for each of the other portions 502-503.

[045] The lower portion of Figure 5 illustrates the manner in which the buffer 500 is indexed. In the Figure 5 embodiment, to access the alternative instruction sequences for both the taken and not taken cases for each of the following branches, the address of a given following branch is used to index the buffer 500. It should be noted that the alternative instruction sequences are stored within the portions 501-503 in an orthogonal manner. In other words, the alternative instruction sequences that can both possibly be taken from a given branch do not reside within the same portion. For example, as depicted in Figure 5, the alternative instruction sequences for branch B1 and B5 can reside within the portion 501 because either the instruction sequence for branch B1 or branch B5 will occur. This is because branch B0 will either be taken or not taken. Thus there is no scenario in which instructions from both branch B1 and branch B5 will occur. Similarly, at the next level removed from branch B0, the alternative instruction sequences for branches B2, B4 and B6 can be stored within the portion 502. These alternative instruction sequences are mutually exclusive in that only one of the three can possibly be executed. Similarly, at the next level, portion 503 stores alternative instruction sequences for the branches B3, B7 and B8.

[046] Figure 6 shows a diagram of an instruction sequence buffer 600 that is used to store instruction sequences for reliably predictable branches that are frequently hit in accordance with one embodiment of the present invention. As illustrated in Figure 6, the buffer 600 comprises four portions 601-604. Each of the portions 601-604 is coupled to respective compare logic components 611-614.

[047] Figure 6 illustrates an alternative use of the storage resources of the instruction sequence buffer. In the Figure 6 embodiment, instead of storing alternative instruction sequences for frequently hit in frequently miss predicted branches, the storage resources are used to store instruction sequences for frequently hit and reliably predicted branches. Thus, instead of storing alternative instruction sequences for the taken case and the not taken case, the storage resources of the buffer 600 are used to store the instruction

sequences of a frequently hit and reliably predictable branch and a number of the subsequent following branches.

[048] Storing the instruction sequences of a frequently hit and reliably predictable branch and a number of the subsequent following branches provides a number of advantages. On subsequent hits due the reliably predictable branch, the instruction sequence of the predictable branch and a number of the subsequent following branches can be provided directly from the buffer 600. This saves a number of cycles with respect to the latency of assembling the reliably predictable branch instruction sequence from the normal decoder module pipeline. In this manner, embodiments of the present invention take advantage of those reliably predictable instruction sequences to cut out latency cycles by providing such sequences directly from the buffer 600.

[049] It should be noted that the buffer 600 is essentially the same structure as the buffer 500 of Figure 5. The difference is in the manner in which the buffer 600 is indexed. As described above, the buffer 600 is used to store reliably predictable instruction sequences that flow from multiple branches. The reliably predict the instruction sequences are stored in multiple ways, shown as way 1 “W1” and way 2 “W2” at the top of each of the portions 601-604. In one embodiment, the address of the branches (e.g., branch B1) are used to index into the cache. For example, in a scenario where a reliably predictable instruction sequence flows from B0 to B1 to B2 to B3, the address of the first following branch B1 is used to index the buffer 600 with the following branches B2 and B3 being used as tags. The following branches B2 and B3 would allow for the same index to be accessed via two different ways with two different tags (b2 and b3). In one embodiment, bits of the branch prediction (e.g., provided from the branch which in table 302) can also be used as tags. In one embodiment, a hash of the following branch B1 and its respective following branches B2 and B3 could be used to access the buffer 600.

[050] The compare logic components 611-614 functions by comparing branch sequence predictions. The components 611-614 compare predictions with sequence hits to score the relative merit of the reliably predictable sequences. For example, if a reliably predictable sequence becomes not so strongly predictable for some reason, this component

will cause it to be evicted from the buffer 600. In one embodiment, if the reliably predict will sequence becomes a frequently hit in frequently miss predicted sequence, the sequence is moved from the accessing and storing methodology illustrated in Figure 6 to the accessing and storing methodology illustrated in Figure 5.

[051] It should be noted that the buffer 600 can be implemented as a unified cache architecture. In such an embodiment, instruction sequences for both the frequently hit in frequently miss predicted instruction sequences and the frequently hit and reliably predictable instruction sequences can both be stored within the common structure of buffer 600. The difference would be the methodology in which they are accessed, indexed and retrieved. In such an embodiment, logic would need to be included to ensure that any instruction sequences that become corrupted due to conflicts or the like are evicted.

[052] Alternatively, in one embodiment, the buffer 600 can be partitioned or otherwise allocated such that storage resources are dedicated to frequently hit frequently miss predicted instruction sequences and dedicated to reliably predictable instruction sequences respectively. Such dedicated allocation would simplify the management of the buffer 600 by making it less likely instruction sequences could corrupt one another.

[053] Figure 7 shows an overview flowchart of the steps of a process 700 for outputting reliably predictable instruction sequences in accordance with one embodiment of the present invention. Process 700 shows exemplary operating steps of, for example, an instruction fetch module of a microprocessor.

[054] Process 700 begins in step 701, where accesses are tracked in order to determine a set of frequently hit instruction sequences. As described above, the sequence predictor assembles predicted instruction sequences. It can determine which instruction sequences are frequently fetched.

[055] In step 702, out of that set, the sequence predictor identifies a branch instruction having a series of subsequent following branch instructions that form a reliably predictable instruction sequence.

[056] In step 703, a reliably predictable instruction sequence for the branch instruction is stored into a buffer (e.g., buffer 600).

[057] In step 704, on a subsequent hit to the branch instruction, the reliably predictable instruction sequence is output from the buffer. This saves a number of cycles with respect to the latency of assembling the reliably predictable branch instruction sequence from the normal decoder module pipeline. In this manner, embodiments of the present invention take advantage of those reliably predictable instruction sequences to cut out latency cycles by providing such sequences directly from the buffer 600.

[058] Figure 8 shows a diagram of an exemplary microprocessor pipeline 800 in accordance with one embodiment of the present invention. The microprocessor pipeline 800 includes a fetch module 801 that implements the functionality of the process for identifying and extracting the instructions comprising an execution, as described above. In the Figure 8 embodiment, the fetch module is followed by a decode module 802, an allocation module 803, a dispatch module 804, an execution module 805 and a retirement modules 806. It should be noted that the microprocessor pipeline 800 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.

[059] The foregoing description, for the purpose of explanation, has been described with reference to specific embodiments. However, the illustrated discussions above are not intended to be exhaustive or to limit the invention to the precise forms disclosed. Many modifications and variations are possible in view of the above teachings. Embodiments were chosen and described in order to best explain the principles of the invention and its practical applications, to thereby enable others skilled in the art to best utilize the invention and various embodiments with various modifications as may be suited to the particular use contemplated.

CLAIMS

What is claimed is:

- 5 1. A method for outputting alternative instruction sequences, comprising
tracking repetitive hits to determine a set of frequently hit instruction sequences for
a microprocessor;
identifying a frequently miss-predicted branch instruction, wherein the predicted
outcome of the branch instruction is frequently wrong;
10 store an alternative instruction sequence for the branch instruction into a buffer;
on a subsequent hit to the branch instruction where the predicted outcome of the
branch instruction was wrong, outputting the alternative instruction sequence from the
buffer.
- 15 2. The method of claim 1, wherein a plurality of miss-predicted branch target
instructions are identified, and a corresponding plurality of respective alternative
instruction sequences are stored in the buffer.
3. The method of claim 1, wherein alternative instruction sequences are stored a
20 number of levels of following branches deep.
4. The method of claim 3, wherein alternative instruction sequences are stored for
taken and not taken conditions of the following branches.
- 25 5. The method of claim 1, wherein outputting the alternative instruction sequence
from the buffer avoids causing a full flush of an instruction pipeline of the microprocessor.
6. The method of claim 1, wherein outputting the alternative instruction sequence
from the buffer reduces a performance penalty from wrongly predicting a branch
30 instruction.

7. The method of claim 1, wherein a tag structure is used to identify the set of frequently hit instruction sequences.

8. A system for outputting alternative instruction sequences, said system

5 comprising:

a fetch module that accesses a plurality of instructions that comprise multiple branch instructions;

a buffer that stores alternative instruction sequences;

10 wherein the fetch module tracks repetitive hits to determine a set of frequently hit instruction sequences for a microprocessor;

where in the fetch module identifies a frequently miss-predicted branch instruction, wherein the predicted outcome of the branch instruction is frequently wrong;

the fetch module stores an alternative instruction sequence for the branch instruction into a buffer;

15 on a subsequent hit to the branch instruction where the predicted outcome of the branch instruction was wrong, the fetch module outputs the alternative instruction sequence from the buffer.

9. The system of claim 8, wherein a plurality of miss-predicted branch instructions
20 are identified, and a corresponding plurality of respective alternative instruction sequences are stored in the buffer.

10. The system of claim 8, wherein alternative instruction sequences are stored a number of levels of following branches deep.

25

11. The system of claim 10, wherein alternative instruction sequences are stored for taken and not taken conditions of the following branches.

12. The system of claim 8, wherein outputting the alternative instruction sequence
30 from the buffer avoids causing a full flush of an instruction pipeline of the microprocessor.

13. The system of claim 8, wherein outputting the alternative instruction sequence from the buffer reduces a performance penalty from wrongly predicting a branch instruction.

5 14. The system of claim 8, wherein a tag structure is used to identify the set of frequently hit instruction sequences.

15. A microprocessor that implements a method of identifying instructions, said microprocessor comprises:

10 a microprocessor pipeline;
 a fetch module included in the microprocessor pipe, wherein the fetch module; and
 a buffer coupled to the fetch module;
 wherein the fetch module tracks repetitive hits to determine a set of frequently hit instruction sequences for a microprocessor;
15 where in the fetch module identifies a frequently miss-predicted branch instruction, wherein the predicted outcome of the branch instruction is frequently wrong;
 the fetch module stores an alternative instruction sequence for the branch instruction into a buffer;
 on a subsequent hit to the branch instruction where the predicted outcome of the
20 branch instruction was wrong, the fetch module outputs the alternative instruction sequence from the buffer.

16. The microprocessor of claim 15, wherein a plurality of miss-predicted branch instructions are identified, and a corresponding plurality of respective alternative
25 instruction sequences are stored in the buffer.

17. The microprocessor of claim 15, wherein alternative instruction sequences are stored a number of levels of following branches deep, and wherein alternative instruction sequences are stored for taken and not taken conditions of the following branches.

30

18. The microprocessor of claim 15, wherein outputting the alternative instruction sequence from the buffer avoids causing a flush of an instruction pipeline of the microprocessor.

5 19. The microprocessor of claim 15, wherein outputting the alternative instruction sequence from the buffer reduces a performance penalty from wrongly predicting a branch instruction.

10 20. The microprocessor of claim 15, wherein a tag structure is used to identify the set of frequently hit instruction sequences.

1/8


100

FIGURE 1

2/8


100

FIGURE 2




FIGURE 3

4/8


400

FIGURE 4

500




FIGURE 5




FIGURE 6

7/8


400

FIGURE 7

8/8




FIGURE 8