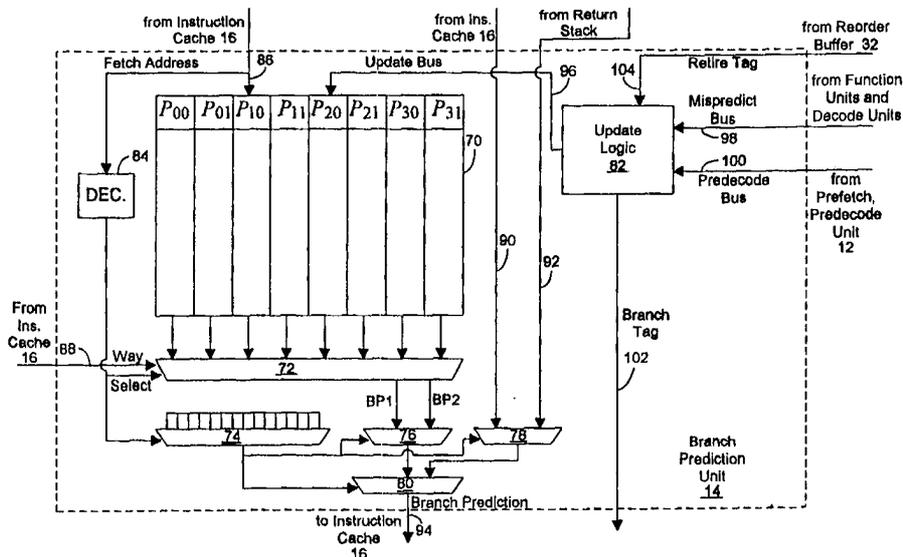




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<p>(51) International Patent Classification <sup>6</sup> : <b>G06F 9/38</b></p>	<p><b>A1</b></p>	<p>(11) International Publication Number: <b>WO 98/22873</b> (43) International Publication Date: 28 May 1998 (28.05.98)</p>
<p>(21) International Application Number: PCT/US97/21048 (22) International Filing Date: 19 November 1997 (19.11.97) (30) Priority Data: 08/752,691 19 November 1996 (19.11.96) US (71) Applicant: ADVANCED MICRO DEVICES, INC. [US/US]; One AMD Place, P.O. Box 3453, Sunnyvale, CA 94088 (US). (72) Inventor: TRAN, Thang, M.; 5209 Eagle Trace Trail, Austin, TX 78730 (US). (74) Agent: KIVLIN, B., Noël; Conley, Rose &amp; Tayon, P.C., P.O. Box 3267, Houston, TX 77253-3267 (US).</p>		<p>(81) Designated States: CN, JP, KR, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).  <b>Published</b> <i>With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i></p>

(54) Title: BRANCH PREDICTION MECHANISM EMPLOYING BRANCH SELECTORS TO SELECT A BRANCH PREDICTION



(57) Abstract

A branch prediction apparatus is provided which stores multiple branch selectors corresponding to instruction bytes within a cache line of instructions or portion thereof. The branch selectors identify a branch prediction to be selected if the corresponding instruction byte is the byte indicated by the offset of the fetch address used to fetch the cache line. Instead of comparing pointers to the branch instructions with the offset of the fetch address, the branch prediction is selected simply by decoding the offset of the fetch address and choosing the corresponding branch selector. The branch prediction apparatus may operate at higher frequencies (i.e. lower clock cycles) than if the pointers to the branch instruction and the fetch address were compared (a greater than or less than comparison). The branch selectors directly determine which branch prediction is appropriate according to the instructions being fetched, thereby decreasing the amount of logic employed to select the branch prediction.

**FOR THE PURPOSES OF INFORMATION ONLY**

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

<b>AL</b>	Albania	<b>ES</b>	Spain	<b>LS</b>	Lesotho	<b>SI</b>	Slovenia
<b>AM</b>	Armenia	<b>FI</b>	Finland	<b>LT</b>	Lithuania	<b>SK</b>	Slovakia
<b>AT</b>	Austria	<b>FR</b>	France	<b>LU</b>	Luxembourg	<b>SN</b>	Senegal
<b>AU</b>	Australia	<b>GA</b>	Gabon	<b>LV</b>	Latvia	<b>SZ</b>	Swaziland
<b>AZ</b>	Azerbaijan	<b>GB</b>	United Kingdom	<b>MC</b>	Monaco	<b>TD</b>	Chad
<b>BA</b>	Bosnia and Herzegovina	<b>GE</b>	Georgia	<b>MD</b>	Republic of Moldova	<b>TG</b>	Togo
<b>BB</b>	Barbados	<b>GH</b>	Ghana	<b>MG</b>	Madagascar	<b>TJ</b>	Tajikistan
<b>BE</b>	Belgium	<b>GN</b>	Guinea	<b>MK</b>	The former Yugoslav Republic of Macedonia	<b>TM</b>	Turkmenistan
<b>BF</b>	Burkina Faso	<b>GR</b>	Greece	<b>ML</b>	Mali	<b>TR</b>	Turkey
<b>BG</b>	Bulgaria	<b>HU</b>	Hungary	<b>MN</b>	Mongolia	<b>TT</b>	Trinidad and Tobago
<b>BJ</b>	Benin	<b>IE</b>	Ireland	<b>MR</b>	Mauritania	<b>UA</b>	Ukraine
<b>BR</b>	Brazil	<b>IL</b>	Israel	<b>MW</b>	Malawi	<b>UG</b>	Uganda
<b>BY</b>	Belarus	<b>IS</b>	Iceland	<b>MX</b>	Mexico	<b>US</b>	United States of America
<b>CA</b>	Canada	<b>IT</b>	Italy	<b>NE</b>	Niger	<b>UZ</b>	Uzbekistan
<b>CF</b>	Central African Republic	<b>JP</b>	Japan	<b>NL</b>	Netherlands	<b>VN</b>	Viet Nam
<b>CG</b>	Congo	<b>KE</b>	Kenya	<b>NO</b>	Norway	<b>YU</b>	Yugoslavia
<b>CH</b>	Switzerland	<b>KG</b>	Kyrgyzstan	<b>NZ</b>	New Zealand	<b>ZW</b>	Zimbabwe
<b>CI</b>	Côte d'Ivoire	<b>KP</b>	Democratic People's Republic of Korea	<b>PL</b>	Poland		
<b>CM</b>	Cameroon	<b>KR</b>	Republic of Korea	<b>PT</b>	Portugal		
<b>CN</b>	China	<b>KZ</b>	Kazakstan	<b>RO</b>	Romania		
<b>CU</b>	Cuba	<b>LC</b>	Saint Lucia	<b>RU</b>	Russian Federation		
<b>CZ</b>	Czech Republic	<b>LI</b>	Liechtenstein	<b>SD</b>	Sudan		
<b>DE</b>	Germany	<b>LK</b>	Sri Lanka	<b>SE</b>	Sweden		
<b>DK</b>	Denmark	<b>LR</b>	Liberia	<b>SG</b>	Singapore		
<b>EE</b>	Estonia						

**TITLE: BRANCH PREDICTION MECHANISM EMPLOYING BRANCH SELECTORS  
TO SELECT A BRANCH PREDICTION**

5

**BACKGROUND OF THE INVENTION**

1. Field of the Invention

This invention relates to the field of microprocessors and, more particularly, to branch prediction mechanisms within microprocessors.

2. Description of the Related Art

Superscalar microprocessors achieve high performance by executing multiple instructions per clock cycle and by choosing the shortest possible clock cycle consistent with the design. As used herein, the term "clock cycle" refers to an interval of time accorded to various stages of an instruction processing pipeline within the microprocessor. Storage devices (e.g. registers and arrays) capture their values according to the clock cycle. For example, a storage device may capture a value according to a rising or falling edge of a clock signal defining the clock cycle. The storage device then stores the value until the subsequent rising or falling edge of the clock signal, respectively. The term "instruction processing pipeline" is used herein to refer to the logic circuits employed to process instructions in a pipelined fashion. Although the pipeline may be divided into any number of stages at which portions of instruction processing are performed, instruction processing generally comprises fetching the instruction, decoding the instruction, executing the instruction, and storing the execution results in the destination identified by the instruction.

An important feature of a superscalar microprocessor (and a superpipelined microprocessor as well) is its branch prediction mechanism. The branch prediction mechanism indicates a predicted direction (taken or not-taken) for a branch instruction, allowing subsequent instruction fetching to continue within the predicted instruction stream indicated by the branch prediction. A branch instruction is an instruction which causes subsequent instructions to be fetched from one of at least two addresses: a sequential address identifying an instruction stream beginning with instructions which directly follow the branch instruction; and a target address identifying an instruction stream beginning at an arbitrary location in memory. Unconditional branch instructions always branch to the target address, while conditional branch instructions may select either the sequential or the target address based on the outcome of a prior instruction. Instructions from the predicted instruction stream may be speculatively executed prior to execution of the branch instruction, and in any case are placed into the instruction processing pipeline prior to execution of the branch instruction. If the predicted instruction stream is correct, then the number of instructions executed per clock cycle is advantageously increased. However, if the predicted instruction stream is incorrect (i.e. one or more branch instructions are predicted incorrectly), then the instructions from the incorrectly predicted instruction stream are discarded from the instruction processing pipeline and the number of instructions executed per clock cycle is decreased.

40

In order to be effective, the branch prediction mechanism must be highly accurate such that the predicted instruction stream is correct as often as possible. Typically, increasing the accuracy of the branch prediction mechanism is achieved by increasing the complexity of the branch prediction mechanism. For example, a cache-line based branch prediction scheme may be employed in which branch predictions are stored with a particular cache line of instruction bytes in an instruction cache. A cache line is a number of contiguous bytes which are treated as a unit for allocation and deallocation of storage space within the instruction cache. When the cache line is fetched, the corresponding branch predictions are also fetched. Furthermore, when the particular cache line is discarded, the corresponding branch predictions are discarded as well. The cache line is aligned in memory. A cache-line based branch prediction scheme may be made more accurate by storing a larger number of branch predictions for each cache line. A given cache line may include multiple branch instructions, each of which is represented by a different branch prediction. Therefore, more branch predictions allocated to a cache line allows for more branch instructions to be represented and predicted by the branch prediction mechanism. A branch instruction which cannot be represented within the branch prediction mechanism is not predicted, and subsequently a "misprediction" may be detected if the branch is found to be taken. However, complexity of the branch prediction mechanism is increased by the need to select between additional branch predictions. As used herein, a "branch prediction" is a value which may be interpreted by the branch prediction mechanism as a prediction of whether or not a branch instruction is taken or not taken. Furthermore, a branch prediction may include the target address. For cache-line based branch prediction mechanisms, a prediction of a sequential line to the cache line being fetched is a branch prediction when no branch instructions are within the instructions being fetched from the cache line.

A problem related to increasing the complexity of the branch prediction mechanism is that the increased complexity generally requires an increased amount of time to form the branch prediction. For example, selecting among multiple branch predictions may require a substantial amount of time. The offset of the fetch address identifies the first byte being fetched within the cache line: a branch prediction for a branch instruction prior to the offset should not be selected. The offset of the fetch address within the cache line may need to be compared to the offset of the branch instructions represented by the branch predictions stored for the cache line in order to determine which branch prediction to use. The branch prediction corresponding to a branch instruction subsequent to the fetch address offset and nearer to the fetch address offset than other branch instructions which are subsequent to the fetch address offset should be selected. As the number of branch predictions is increased, the complexity (and time required) for the selection logic increases. When the amount of time needed to form a branch prediction for a fetch address exceeds the clock cycle time of the microprocessor, performance of the microprocessor may be decreased. Because the branch prediction cannot be formed in a single clock cycle, "bubbles" are introduced into the instruction processing pipeline during clock cycles that instructions cannot be fetched due to a lack of a branch prediction corresponding to a previous fetch address. The bubble occupies various stages in the instruction processing pipeline during subsequent clock cycles, and no work occurs at the stage including the bubble because no instructions are included in the bubble. Performance of the microprocessor may thereby be decreased.

SUMMARY OF THE INVENTION

The problems outlined above are in large part solved by a branch prediction apparatus in accordance with the present invention. The branch prediction apparatus stores multiple branch selectors corresponding to instruction bytes within a cache line of instructions or portion thereof. The branch selectors identify a branch prediction to be selected if the corresponding instruction byte is the byte indicated by the offset of the fetch address used to fetch the cache line. Instead of comparing pointers to the branch instructions with the offset of the fetch address, the branch prediction is selected simply by decoding the offset of the fetch address and choosing the corresponding branch selector. Advantageously, the branch prediction apparatus may operate at a higher frequencies (i.e. lower clock cycles) than if the pointers to the branch instruction and the fetch address were compared (a greater than or less than comparison). The branch selectors directly determine which branch prediction is appropriate according to the instructions being fetched, thereby decreasing the amount of logic employed to select the branch prediction.

Broadly speaking, the present invention contemplates a method for selecting a branch prediction corresponding to a group of contiguous instruction bytes including a plurality of instructions. A plurality of branch selectors are stored in a branch prediction storage, wherein at least one of the plurality of branch selectors corresponds to a first one of the plurality of instructions. The branch selector identifies a particular branch prediction to be selected if the first one of the plurality of instructions is fetched. The group of contiguous instruction bytes is fetched concurrent with fetching the plurality of branch selectors. The fetch address identifies the group of contiguous instruction bytes. One of the plurality of branch selectors is selected in response to the fetch address. The branch prediction is selected in response to the one of the plurality of the branch selectors.

The present invention further contemplates a branch prediction apparatus, comprising a branch prediction storage and a selection mechanism. The branch prediction storage is coupled to receive a fetch address corresponding to a group of contiguous instruction bytes being fetched from an instruction cache. The branch prediction storage is configured to store a plurality of branch selectors wherein at least one of the plurality of branch selectors corresponds to a first instruction within the group of contiguous instruction bytes. The at least one of the plurality of branch selectors identifies a particular branch prediction to be selected if the first instruction is fetched. Coupled to the branch prediction storage to receives the plurality of branch selectors, the selection mechanism is configured to select a particular one of the plurality of branch selectors in response to a plurality of least significant bits of a fetch address used to fetch the group of contiguous instruction bytes.

The present invention still further contemplates a microprocessor comprising an instruction cache and a branch prediction unit. The instruction cache is configured to store a plurality of cache lines of instruction bytes and to provide a group of instruction bytes upon receipt of a fetch address to an instruction processing pipeline of the microprocessor. Coupled to the instruction cache and coupled to receive the fetch address concurrent with the instruction cache, the branch prediction unit is configured to store a plurality of branch selectors with respect to the group of instruction bytes and is configured to select one of the plurality of

branch selectors in response to the fetch address. The one of the plurality of branch selectors identifies a branch prediction which is used as a subsequent fetch address by the instruction cache.

### BRIEF DESCRIPTION OF THE DRAWINGS

5

Other objects and advantages of the invention will become apparent upon reading the following detailed description and upon reference to the accompanying drawings in which:

Fig. 1 is a block diagram of one embodiment of a superscalar microprocessor.

10

Fig. 2 is a block diagram of one embodiment of a pair of decode units shown in Fig. 1.

Fig. 3 is a diagram of a group of contiguous instruction bytes and a corresponding set of branch selectors.

15

Fig. 4 is a block diagram of a portion of one embodiment of a branch prediction unit shown in Fig. 1.

Fig. 5 is a diagram of a prediction block for a group of contiguous instruction bytes as stored in the branch prediction unit shown in Fig. 4.

20

Fig. 6 is a table showing an exemplary encoding of a branch selector.

Fig. 7 is a flowchart depicting steps performed in order to update a set of branch selectors corresponding to a group of contiguous instruction bytes.

25

Fig. 8 is a first example of updating the set of branch selectors.

Fig. 9 is a second example of updating the set of branch selectors.

30

Fig. 10 is a third example of updating the set of branch selectors.

Fig. 11 is a block diagram of a computer system including the microprocessor shown in Fig. 1.

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

**DETAILED DESCRIPTION OF THE INVENTION**

Turning now to Fig. 1, a block diagram of one embodiment of a microprocessor 10 is shown.

Microprocessor 10 includes a prefetch/predecode unit 12, a branch prediction unit 14, an instruction cache 16,  
5 an instruction alignment unit 18, a plurality of decode units 20A-20C, a plurality of reservation stations 22A-  
22C, a plurality of functional units 24A-24C, a load/store unit 26, a data cache 28, a register file 30, a reorder  
buffer 32, and an MROM unit 34. Elements referred to herein with a particular reference number followed by  
a letter will be collectively referred to by the reference number alone. For example, decode units 20A-20C  
will be collectively referred to as decode units 20.

10 Prefetch/predecode unit 12 is coupled to receive instructions from a main memory subsystem (not  
shown), and is further coupled to instruction cache 16 and branch prediction unit 14. Similarly, branch  
prediction unit 14 is coupled to instruction cache 16. Still further, branch prediction unit 14 is coupled to  
decode units 20 and functional units 24. Instruction cache 16 is further coupled to MROM unit 34 and  
instruction alignment unit 18. Instruction alignment unit 18 is in turn coupled to decode units 20. Each  
15 decode unit 20A-20C is coupled to load/store unit 26 and to respective reservation stations 22A-22C.  
Reservation stations 22A-22C are further coupled to respective functional units 24A-24C. Additionally,  
decode units 20 and reservation stations 22 are coupled to register file 30 and reorder buffer 32. Functional  
units 24 are coupled to load/store unit 26, register file 30, and reorder buffer 32 as well. Data cache 28 is  
coupled to load/store unit 26 and to the main memory subsystem. Finally, MROM unit 34 is coupled to decode  
20 units 20.

Generally speaking, branch prediction unit 14 employs a cache-line based branch prediction  
mechanism for predicting branch instructions. Multiple branch predictions may be stored for each cache line.  
Additionally, a branch selector is stored for each byte within the cache line. The branch selector for a  
particular byte indicates which of the branch predictions which may be stored with respect to the cache line is  
25 the branch prediction appropriate for an instruction fetch address which fetches the particular byte. The  
appropriate branch prediction is the branch prediction for the first predicted-taken branch instruction  
encountered within the cache line subsequent to the particular byte. As used herein, the terms "subsequent"  
and "prior to" refer to an ordering of bytes within the cache line. A byte stored at a memory address which is  
numerically smaller than the memory address at which a second byte is stored is prior to the second byte.  
30 Conversely, a byte stored at a memory address which is numerically larger than the memory address of a  
second byte is subsequent to the second byte. Similarly, a first instruction is prior to a second instruction in  
program order if the first instruction is encountered before the second instruction when stepping one at a time  
through the sequence of instructions forming the program.

In one embodiment, microprocessor 10 employs a microprocessor architecture in which the  
35 instruction set is a variable byte length instruction set (e.g. the x86 microprocessor architecture). When a  
variable byte length instruction set is employed, any byte within the cache line may be identified as the first  
byte to be fetched by a given fetch address. For example, a branch instruction may have a target address at  
byte position two within a cache line. In such a case, the bytes at byte positions zero and one are not being  
fetched during the current cache access. Additionally, bytes subsequent to a predicted-taken branch which is

subsequent to the first byte are not fetched during the current cache access. Since branch selectors are stored for each byte, the branch prediction for the predicted taken branch can be located by selecting the branch selector of the first byte to be fetched from the cache line. The branch selector is used to select the appropriate branch prediction, which is then provided to the instruction fetch logic in instruction cache 16. During the succeeding clock cycle, the branch prediction is used as the fetch address. Advantageously, the process of comparing the byte position of the first byte being fetched to the byte positions of the predicted-taken branch instructions is eliminated from the generation of a branch prediction in response to a fetch address. The amount of time required to form a branch prediction may be reduced accordingly, allowing the branch prediction mechanism to operate at higher clock frequencies (i.e. shorter clock cycles) while still providing a single cycle branch prediction.

It is noted that, although the term "cache line" has been used in the preceding discussion, some embodiments of instruction cache 16 may not provide an entire cache line at its output during a given clock cycle. For example, in one embodiment instruction cache 16 is configured with 32 byte cache lines. However, only 16 bytes are fetched in a given clock cycle (either the upper half or the lower half of the cache line). The branch prediction storage locations and branch selectors are allocated to the portion of the cache line being fetched. As used herein, the term "group of contiguous instruction bytes" is used to refer to the instruction bytes which are provided by the instruction cache in a particular clock cycle in response to a fetch address. A group of contiguous instruction bytes may be a portion of a cache line or an entire cache line, according to various embodiments. When a group of contiguous instruction bytes is a portion of a cache line, it is still an aligned portion of a cache line. For example, if a group of contiguous instruction bytes is half a cache line, it is either the upper half of the cache line or the lower half of the cache line. A number of branch prediction storage locations are allocated to each group of contiguous instruction bytes, and branch selectors indicate one of the branch prediction storage locations associated with that group. Furthermore, branch selectors may indicate a return stack address from a return stack structure or a sequential address if no branch instructions are encountered between the corresponding byte and the last byte in the group of contiguous instruction bytes.

Instruction cache 16 is a high speed cache memory provided to store instructions. Instructions are fetched from instruction cache 16 and dispatched to decode units 20. In one embodiment, instruction cache 16 is configured to store up to 32 kilobytes of instructions in a 4 way set associative structure having 32 byte lines (a byte comprises 8 binary bits). Instruction cache 16 may additionally employ a way prediction scheme in order to speed access times to the instruction cache. Instead of accessing tags identifying each line of instructions and comparing the tags to the fetch address to select a way, instruction cache 16 predicts the way that is accessed. In this manner, the way is selected prior to accessing the instruction storage. The access time of instruction cache 16 may be similar to a direct-mapped cache. A tag comparison is performed and, if the way prediction is incorrect, the correct instructions are fetched and the incorrect instructions are discarded. It is noted that instruction cache 16 may be implemented as a fully associative, set associative, or direct mapped configuration.

Instructions are fetched from main memory and stored into instruction cache 16 by prefetch/predecode unit 12. Instructions may be prefetched prior to the request thereof from instruction cache 16 in accordance with a prefetch scheme. A variety of prefetch schemes may be employed by

prefetch/predecode unit 12. As prefetch/predecode unit 12 transfers instructions from main memory to instruction cache 16, prefetch/predecode unit 12 generates three predecode bits for each byte of the instructions: a start bit, an end bit, and a functional bit. The predecode bits form tags indicative of the boundaries of each instruction. The predecode tags may also convey additional information such as whether a given instruction can be decoded directly by decode units 20 or whether the instruction is executed by invoking a microcode procedure controlled by MROM unit 34, as will be described in greater detail below. Still further, prefetch/predecode unit 12 may be configured to detect branch instructions and to store branch prediction information corresponding to the branch instructions into branch prediction unit 14.

One encoding of the predecode tags for an embodiment of microprocessor 10 employing the x86 instruction set will next be described. If a given byte is the first byte of an instruction, the start bit for that byte is set. If the byte is the last byte of an instruction, the end bit for that byte is set. Instructions which may be directly decoded by decode units 20 are referred to as "fast path" instructions. The remaining x86 instructions are referred to as MROM instructions, according to one embodiment. For fast path instructions, the functional bit is set for each prefix byte included in the instruction, and cleared for other bytes. Alternatively, for MROM instructions, the functional bit is cleared for each prefix byte and set for other bytes. The type of instruction may be determined by examining the functional bit corresponding to the end byte. If that functional bit is clear, the instruction is a fast path instruction. Conversely, if that functional bit is set, the instruction is an MROM instruction. The opcode of an instruction may thereby be located within an instruction which may be directly decoded by decode units 20 as the byte associated with the first clear functional bit in the instruction. For example, a fast path instruction including two prefix bytes, a Mod R/M byte, and an SIB byte would have start, end, and functional bits as follows:

Start bits	10000
End bits	00001
Functional bits	11000

MROM instructions are instructions which are determined to be too complex for decode by decode units 20. MROM instructions are executed by invoking MROM unit 34. More specifically, when an MROM instruction is encountered, MROM unit 34 parses and issues the instruction into a subset of defined fast path instructions to effectuate the desired operation. MROM unit 34 dispatches the subset of fast path instructions to decode units 20. A listing of exemplary x86 instructions categorized as fast path instructions will be provided further below.

Microprocessor 10 employs branch prediction in order to speculatively fetch instructions subsequent to conditional branch instructions. Branch prediction unit 14 is included to perform branch prediction operations. In one embodiment, up to two branch target addresses are stored with respect to each 16 byte portion of each cache line in instruction cache 16. Prefetch/predecode unit 12 determines initial branch targets when a particular line is predecoded. Subsequent updates to the branch targets corresponding to a cache line may occur due to the execution of instructions within the cache line. Instruction cache 16 provides an indication of the instruction address being fetched, so that branch prediction unit 14 may determine which branch target addresses to select for forming a branch prediction. Decode units 20 and functional units 24

provide update information to branch prediction unit 14. Because branch prediction unit 14 stores two targets per 16 byte portion of the cache line, some branch instructions within the line may not be stored in branch prediction unit 14. Decode units 20 detect branch instructions which were not predicted by branch prediction unit 14. Functional units 24 execute the branch instructions and determine if the predicted branch direction is incorrect. The branch direction may be "taken", in which subsequent instructions are fetched from the target address of the branch instruction. Conversely, the branch direction may be "not taken", in which subsequent instructions are fetched from memory locations consecutive to the branch instruction. When a mispredicted branch instruction is detected, instructions subsequent to the mispredicted branch are discarded from the various units of microprocessor 10. A variety of suitable branch prediction algorithms may be employed by branch prediction unit 14.

Instructions fetched from instruction cache 16 are conveyed to instruction alignment unit 18. As instructions are fetched from instruction cache 16, the corresponding predecode data is scanned to provide information to instruction alignment unit 18 (and to MROM unit 34) regarding the instructions being fetched. Instruction alignment unit 18 utilizes the scanning data to align an instruction to each of decode units 20. In one embodiment, instruction alignment unit 18 aligns instructions from three sets of eight instruction bytes to decode units 20. Instructions are selected independently from each set of eight instruction bytes into preliminary issue positions. The preliminary issue positions are then merged to a set of aligned issue positions corresponding to decode units 20, such that the aligned issue positions contain the three instructions which are prior to other instructions within the preliminary issue positions in program order. Decode unit 20A receives an instruction which is prior to instructions concurrently received by decode units 20B and 20C (in program order). Similarly, decode unit 20B receives an instruction which is prior to the instruction concurrently received by decode unit 20C in program order.

Decode units 20 are configured to decode instructions received from instruction alignment unit 18. Register operand information is detected and routed to register file 30 and reorder buffer 32. Additionally, if the instructions require one or more memory operations to be performed, decode units 20 dispatch the memory operations to load/store unit 26. Each instruction is decoded into a set of control values for functional units 24, and these control values are dispatched to reservation stations 22 along with operand address information and displacement or immediate data which may be included with the instruction.

Microprocessor 10 supports out of order execution, and thus employs reorder buffer 32 to keep track of the original program sequence for register read and write operations, to implement register renaming, to allow for speculative instruction execution and branch misprediction recovery, and to facilitate precise exceptions. A temporary storage location within reorder buffer 32 is reserved upon decode of an instruction that involves the update of a register to thereby store speculative register states. If a branch prediction is incorrect, the results of speculatively-executed instructions along the mispredicted path can be invalidated in the buffer before they are written to register file 30. Similarly, if a particular instruction causes an exception, instructions subsequent to the particular instruction may be discarded. In this manner, exceptions are "precise" (i.e. instructions subsequent to the particular instruction causing the exception are not completed prior to the exception). It is noted that a particular instruction is speculatively executed if it is executed prior to instructions which precede the particular instruction in program order. Preceding instructions may be a branch

instruction or an exception-causing instruction, in which case the speculative results may be discarded by reorder buffer 32.

The instruction control values and immediate or displacement data provided at the outputs of decode units 20 are routed directly to respective reservation stations 22. In one embodiment, each reservation station 22 is capable of holding instruction information (i.e., instruction control values as well as operand values, operand tags and/or immediate data) for up to three pending instructions awaiting issue to the corresponding functional unit. It is noted that for the embodiment of Fig. 1, each reservation station 22 is associated with a dedicated functional unit 24. Accordingly, three dedicated "issue positions" are formed by reservation stations 22 and functional units 24. In other words, issue position 0 is formed by reservation station 22A and functional unit 24A. Instructions aligned and dispatched to reservation station 22A are executed by functional unit 24A. Similarly, issue position 1 is formed by reservation station 22B and functional unit 24B; and issue position 2 is formed by reservation station 22C and functional unit 24C.

Upon decode of a particular instruction, if a required operand is a register location, register address information is routed to reorder buffer 32 and register file 30 simultaneously. Those of skill in the art will appreciate that the x86 register file includes eight 32 bit real registers (i.e., typically referred to as EAX, EBX, ECX, EDX, EBP, ESI, EDI and ESP). In embodiments of microprocessor 10 which employ the x86 microprocessor architecture, register file 30 comprises storage locations for each of the 32 bit real registers. Additional storage locations may be included within register file 30 for use by MROM unit 34. Reorder buffer 32 contains temporary storage locations for results which change the contents of these registers to thereby allow out of order execution. A temporary storage location of reorder buffer 32 is reserved for each instruction which, upon decode, is determined to modify the contents of one of the real registers. Therefore, at various points during execution of a particular program, reorder buffer 32 may have one or more locations which contain the speculatively executed contents of a given register. If following decode of a given instruction it is determined that reorder buffer 32 has a previous location or locations assigned to a register used as an operand in the given instruction, the reorder buffer 32 forwards to the corresponding reservation station either: 1) the value in the most recently assigned location, or 2) a tag for the most recently assigned location if the value has not yet been produced by the functional unit that will eventually execute the previous instruction. If reorder buffer 32 has a location reserved for a given register, the operand value (or reorder buffer tag) is provided from reorder buffer 32 rather than from register file 30. If there is no location reserved for a required register in reorder buffer 32, the value is taken directly from register file 30. If the operand corresponds to a memory location, the operand value is provided to the reservation station through load/store unit 26.

In one particular embodiment, reorder buffer 32 is configured to store and manipulate concurrently decoded instructions as a unit. This configuration will be referred to herein as "line-oriented". By manipulating several instructions together, the hardware employed within reorder buffer 32 may be simplified. For example, a line-oriented reorder buffer included in the present embodiment allocates storage sufficient for instruction information pertaining to three instructions (one from each decode unit 20) whenever one or more instructions are dispatched by decode units 20. By contrast, a variable amount of storage is allocated in conventional reorder buffers, dependent upon the number of instructions actually dispatched. A comparatively larger number of logic gates may be required to allocate the variable amount of storage. When each of the

concurrently decoded instructions has executed, the instruction results are stored into register file 30 simultaneously. The storage is then free for allocation to another set of concurrently decoded instructions. Additionally, the amount of control logic circuitry employed per instruction is reduced because the control logic is amortized over several concurrently decoded instructions. A reorder buffer tag identifying a particular instruction may be divided into two fields: a line tag and an offset tag. The line tag identifies the set of concurrently decoded instructions including the particular instruction, and the offset tag identifies which instruction within the set corresponds to the particular instruction. It is noted that storing instruction results into register file 30 and freeing the corresponding storage is referred to as "retiring" the instructions. It is further noted that any reorder buffer configuration may be employed in various embodiments of microprocessor 10.

As noted earlier, reservation stations 22 store instructions until the instructions are executed by the corresponding functional unit 24. An instruction is selected for execution if: (i) the operands of the instruction have been provided; and (ii) the operands have not yet been provided for instructions which are within the same reservation station 22A-22C and which are prior to the instruction in program order. It is noted that when an instruction is executed by one of the functional units 24, the result of that instruction is passed directly to any reservation stations 22 that are waiting for that result at the same time the result is passed to update reorder buffer 32 (this technique is commonly referred to as "result forwarding"). An instruction may be selected for execution and passed to a functional unit 24A-24C during the clock cycle that the associated result is forwarded. Reservation stations 22 route the forwarded result to the functional unit 24 in this case.

In one embodiment, each of the functional units 24 is configured to perform integer arithmetic operations of addition and subtraction, as well as shifts, rotates, logical operations, and branch operations. The operations are performed in response to the control values decoded for a particular instruction by decode units 20. It is noted that a floating point unit (not shown) may also be employed to accommodate floating point operations. The floating point unit may be operated as a coprocessor, receiving instructions from MROM unit 34 and subsequently communicating with reorder buffer 32 to complete the instructions. Additionally, functional units 24 may be configured to perform address generation for load and store memory operations performed by load/store unit 26.

Each of the functional units 24 also provides information regarding the execution of conditional branch instructions to the branch prediction unit 14. If a branch prediction was incorrect, branch prediction unit 14 flushes instructions subsequent to the mispredicted branch that have entered the instruction processing pipeline, and causes fetch of the required instructions from instruction cache 16 or main memory. It is noted that in such situations, results of instructions in the original program sequence which occur after the mispredicted branch instruction are discarded, including those which were speculatively executed and temporarily stored in load/store unit 26 and reorder buffer 32.

Results produced by functional units 24 are sent to reorder buffer 32 if a register value is being updated, and to load/store unit 26 if the contents of a memory location are changed. If the result is to be stored in a register, reorder buffer 32 stores the result in the location reserved for the value of the register when the instruction was decoded. A plurality of result buses 38 are included for forwarding of results from functional units 24 and load/store unit 26. Result buses 38 convey the result generated, as well as the reorder buffer tag

identifying the instruction being executed.

Load/store unit 26 provides an interface between functional units 24 and data cache 28. In one embodiment, load/store unit 26 is configured with a load/store buffer having eight storage locations for data and address information for pending loads or stores. Decode units 20 arbitrate for access to the load/store unit 26. When the buffer is full, a decode unit must wait until load/store unit 26 has room for the pending load or store request information. Load/store unit 26 also performs dependency checking for load memory operations against pending store memory operations to ensure that data coherency is maintained. A memory operation is a transfer of data between microprocessor 10 and the main memory subsystem. Memory operations may be the result of an instruction which utilizes an operand stored in memory, or may be the result of a load/store instruction which causes the data transfer but no other operation. Additionally, load/store unit 26 may include a special register storage for special registers such as the segment registers and other registers related to the address translation mechanism defined by the x86 microprocessor architecture.

In one embodiment, load/store unit 26 is configured to perform load memory operations speculatively. Store memory operations are performed in program order, but may be speculatively stored into the predicted way. If the predicted way is incorrect, the data prior to the store memory operation is subsequently restored to the predicted way and the store memory operation is performed to the correct way. In another embodiment, stores may be executed speculatively as well. Speculatively executed stores are placed into a store buffer, along with a copy of the cache line prior to the update. If the speculatively executed store is later discarded due to branch misprediction or exception, the cache line may be restored to the value stored in the buffer. It is noted that load/store unit 26 may be configured to perform any amount of speculative execution, including no speculative execution.

Data cache 28 is a high speed cache memory provided to temporarily store data being transferred between load/store unit 26 and the main memory subsystem. In one embodiment, data cache 28 has a capacity of storing up to sixteen kilobytes of data in an eight way set associative structure. Similar to instruction cache 16, data cache 28 may employ a way prediction mechanism. It is understood that data cache 28 may be implemented in a variety of specific memory configurations, including a set associative configuration.

In one particular embodiment of microprocessor 10 employing the x86 microprocessor architecture, instruction cache 16 and data cache 28 are linearly addressed. The linear address is formed from the offset specified by the instruction and the base address specified by the segment portion of the x86 address translation mechanism. Linear addresses may optionally be translated to physical addresses for accessing a main memory. The linear to physical translation is specified by the paging portion of the x86 address translation mechanism. It is noted that a linear addressed cache stores linear address tags. A set of physical tags (not shown) may be employed for mapping the linear addresses to physical addresses and for detecting translation aliases. Additionally, the physical tag block may perform linear to physical address translation.

Turning now to Fig. 2, a block diagram of one embodiment of decode units 20B and 20C is shown. Each decode unit 20 receives an instruction from instruction alignment unit 18. Additionally, MROM unit 34 is coupled to each decode unit 20 for dispatching fast path instructions corresponding to a particular MROM instruction. Decode unit 20B comprises early decode unit 40B, multiplexor 42B, and opcode decode unit

44B. Similarly, decode unit 20C includes early decode unit 40C, multiplexor 42C, and opcode decode unit 44C.

Certain instructions in the x86 instruction set are both fairly complicated and frequently used. In one embodiment of microprocessor 10, such instructions include more complex operations than the hardware included within a particular functional unit 24A-24C is configured to perform. Such instructions are classified as a special type of MROM instruction referred to as a "double dispatch" instruction. These instructions are dispatched to a pair of opcode decode units 44. It is noted that opcode decode units 44 are coupled to respective reservation stations 22. Each of opcode decode units 44A-44C forms an issue position with the corresponding reservation station 22A-22C and functional unit 24A-24C. Instructions are passed from an opcode decode unit 44 to the corresponding reservation station 22 and further to the corresponding functional unit 24.

Multiplexor 42B is included for selecting between the instructions provided by MROM unit 34 and by early decode unit 40B. During times in which MROM unit 34 is dispatching instructions, multiplexor 42B selects instructions provided by MROM unit 34. At other times, multiplexor 42B selects instructions provided by early decode unit 40B. Similarly, multiplexor 42C selects between instructions provided by MROM unit 34, early decode unit 40B, and early decode unit 40C. The instruction from MROM unit 34 is selected during times in which MROM unit 34 is dispatching instructions. During times in which the early decode unit within decode unit 20A (not shown) detects a double dispatch instruction, the instruction from early decode unit 40B is selected by multiplexor 42C. Otherwise, the instruction from early decode unit 40C is selected. Selecting the instruction from early decode unit 40B into opcode decode unit 44C allows a fast path instruction decoded by decode unit 20B to be dispatched concurrently with a double dispatch instruction decoded by decode unit 20A.

According to one embodiment employing the x86 instruction set, early decode units 40 perform the following operations:

- (i) merge the prefix bytes of the instruction into an encoded prefix byte;
- (ii) decode unconditional branch instructions (which may include the unconditional jump, the CALL, and the RETURN) which were not detected during branch prediction;
- (iii) decode source and destination flags;
- (iv) decode the source and destination operands which are register operands and generate operand size information; and
- (v) determine the displacement and/or immediate size so that displacement and immediate data may be routed to the opcode decode unit.

Opcode decode units 44 are configured to decode the opcode of the instruction, producing control values for functional unit 24. Displacement and immediate data are routed with the control values to reservation stations 22.

Since early decode units 40 detect operands, the outputs of multiplexors 42 are routed to register file 30 and reorder buffer 32. Operand values or tags may thereby be routed to reservation stations 22. Additionally, memory operands are detected by early decode units 40. Therefore, the outputs of multiplexors

42 are routed to load/store unit 26. Memory operations corresponding to instructions having memory operands are stored by load/store unit 26.

Turning now to Fig. 3, a diagram of an exemplary group of contiguous instruction bytes 50 and a corresponding set of branch selectors 52 are shown. In Fig. 3, each byte within an instruction is illustrated by a short vertical line (e.g. reference number 54). Additionally, the vertical lines separating instructions in group 50 delimit bytes (e.g. reference number 56). The instructions shown in Fig. 3 are variable in length, and therefore the instruction set including the instructions shown in Fig. 3 is a variable byte length instruction set. In other words, a first instruction within the variable byte length instruction set may occupy a first number of bytes which is different than a second number of bytes occupied by a second instruction within the instruction set. Other instruction sets may be fixed-length, such that each instruction within the instruction set occupies the same number of bytes as each other instruction.

As illustrated in Fig. 3, group 50 includes non-branch instructions IN0-IN5. Instructions IN0, IN3, IN4, and IN5 are two byte instructions. Instruction IN1 is a one byte instruction and instruction IN2 is a three byte instruction. Two predicted-taken branch instructions PB0 and PB1 are illustrated as well, each shown as occupying two bytes. It is noted that both non-branch and branch instructions may occupy various numbers of bytes.

The end byte of each predicted-taken branch PB0 and PB1 provides a division of group 50 into three regions: a first region 58, a second region 60, and a third region 62. If a fetch address identifying group 50 is presented, and the offset of the fetch address within the group identifies a byte position within first region 58, then the first predicted-taken branch instruction to be encountered is PB0 and therefore the branch prediction for PB0 is selected by the branch prediction mechanism. Similarly, if the offset of the fetch address identifies a byte within second region 60, the appropriate branch prediction is the branch prediction for PB1. Finally, if the offset of the fetch address identifies a byte within third region 62, then there is no predicted-taken branch instruction within the group of instruction bytes and subsequent to the identified byte. Therefore, the branch prediction for third region 62 is sequential. The sequential address identifies the group of instruction bytes which immediately follows group 50 within main memory.

As used herein, the offset of an address comprises a number of least significant bits of the address. The number is sufficient to provide different encodings of the bits for each byte within the group of bytes to which the offset relates. For example, group 50 is 16 bytes. Therefore, four least significant bits of an address within the group form the offset of the address. The remaining bits of the address identify group 50 from other groups of instruction bytes within the main memory. Additionally, a number of least significant bits of the remaining bits form an index used by instruction cache 16 to select a row of storage locations which are eligible for storing group 50.

Set 52 is an exemplary set of branch selectors for group 50. One branch selector is included for each byte within group 50. The branch selectors within set 52 use the encoding shown in Fig. 6 below. In the example, the branch prediction for PB0 is stored as the second of two branch predictions associated with group 50 (as indicated by a branch selector value of "3"). Therefore, the branch selector for each byte within first region 58 is set to "3". Similarly, the branch prediction for PB1 is stored as the first of the branch predictions (as indicated by a branch selector value of "2"). Therefore, the branch selector for each byte within

second region 60 is set to "2". Finally, the sequential branch prediction is indicated by the branch selectors for bytes within third region 62 by a branch selector encoding of "0".

It is noted that, due to the variable byte length nature of the x86 instruction set, a branch instruction may begin within one group of contiguous instruction bytes and end within a second group of contiguous instruction bytes. In such a case, the branch prediction for the branch instruction is stored with the second group of contiguous instruction bytes. Among other things, the bytes of the branch instruction which are stored within the second group of contiguous instruction bytes need to be fetched and dispatched. Forming the branch prediction in the first group of contiguous instruction bytes would cause the bytes of the branch instruction which lie within the second group of instruction bytes not to be fetched.

Turning now to Fig. 4, a portion of one embodiment of branch prediction unit 14 is shown. Other embodiments of branch prediction unit 14 and the portion shown in Fig. 4 are contemplated. As shown in Fig. 4, branch prediction unit 14 includes a branch prediction storage 70, a way multiplexor 72, a branch selector multiplexor 74, a branch prediction multiplexor 76, a sequential/return multiplexor 78, a final prediction multiplexor 80, an update logic block 82, and a decoder 84. Branch prediction storage 70 and decoder 84 are coupled to a fetch address bus 86 from instruction cache 16. A fetch address concurrently provided to the instruction bytes storage within instruction cache 16 is conveyed upon fetch address bus 86. Decoder block 84 provides selection controls to prediction selector multiplexor 74. Prediction controls for way multiplexor 72 are provided via a way selection bus 88 from instruction cache 16. Way selection bus 88 provides the way of instruction cache 16 which is storing the cache line corresponding to the fetch address provided on fetch address bus 86. Additionally, a selection control is provided by decoder 84 based upon which portion of the cache line is being fetched. Way multiplexor 72 is coupled to receive the contents of each storage location within the row of branch prediction storage 70 which is indexed by the fetch address upon fetch address bus 86. Branch selector multiplexor 74 and branch prediction multiplexor 76 are coupled to receive portions of the output of way multiplexor 72 as inputs. Additionally, the output of branch selector multiplexor 74 provides selection controls for multiplexors 76, 78, and 80. Sequential/return multiplexor 78 selects between a sequential address provided upon a sequential address bus 90 from instruction cache 16 and a return address provided upon a return address bus 92 from a return stack. The output of multiplexors 76 and 78 is provided to final prediction multiplexor 80, which provides a branch prediction bus 94 to instruction cache 16. Instruction cache 16 uses the branch prediction provided upon branch prediction bus 94 as the fetch address for the subsequent clock cycle. Update logic block 82 is coupled to branch prediction storage 70 via an update bus 96 used to update branch prediction information stored therein. Update logic block 82 provides updates in response to a misprediction signalled via a mispredict bus 98 from functional units 24 and decode units 20. Additionally, update logic block 82 provides updates in response to newly predecoded instruction indicated by prefetch/predecode unit 12 upon a predecode bus 100.

Branch prediction storage 70 is arranged with a number of ways equal to the number of ways in instruction cache 16. For each way, a prediction block is stored for each group of contiguous instruction bytes existing within a cache line. In the embodiment of Fig. 4, two groups of instruction bytes are included in each cache line. Therefore, prediction block  $P_{00}$  is the prediction block corresponding to the first group of contiguous instruction bytes in the first way and prediction block  $P_{01}$  is the prediction block corresponding to

the second group of contiguous instruction bytes in the first way. Similarly, prediction block  $P_{10}$  is the prediction block corresponding to the first group of contiguous instruction bytes in the second way and prediction block  $P_{11}$  is the prediction block corresponding to the second group of contiguous instruction bytes in the second way, etc. Each prediction block  $P_{00}$  to  $P_{31}$  in the indexed row is provided as an output of branch prediction storage 70, and hence as an input to way multiplexor 72. The indexed row is similar to indexing into a cache: a number of bits which are not part of the offset portion of the fetch address are used to select one of the rows of branch prediction storage 70. It is noted that branch prediction storage 70 may be configured with fewer rows than instruction cache 16. For example, branch prediction storage 70 may include 1/4 the number of rows of instruction cache 16. In such a case, the address bits which are index bits of instruction cache 16 but which are not index bits of branch prediction storage 70 may be stored with the branch prediction information and checked against the corresponding bits of the fetch address to confirm that the branch prediction information is associated with the row of instruction cache 16 which is being accessed.

Way multiplexor 72 selects one of the sets of branch prediction information  $P_{00}$ - $P_{31}$  based upon the way selection provided from instruction cache 16 and the group of instruction bytes referenced by the fetch address. In the embodiment shown, for example, a 32 byte cache line is divided into two 16 byte groups. Therefore, the fifth least significant bit of the address is used to select which of the two groups contains the fetch address. If the fifth least significant bit is zero, then the first group of contiguous instruction bytes is selected. If the fifth least significant bit is one, then the second group of contiguous instruction bytes is selected. It is noted that the way selection provided upon way select bus 88 may be a way prediction produced by a branch prediction from the previous clock cycle, according to one embodiment. Alternatively, the way selection may be generated via tag comparisons between the fetch address and the address tags identifying the cache lines stored in each way of the instruction cache. It is noted that an address tag is the portion of the address which is not an offset within the cache line nor an index into the instruction cache.

The selected prediction block provided by way multiplexor 72 includes branch selectors for each byte in the group of contiguous instruction bytes, as well as branch predictions BP1 and BP2. The branch selectors are provided to branch selector multiplexor 74, which selects one of the branch selectors based upon selection controls provided by decoder 84. Decoder 84 decodes the offset of the fetch address into the group of contiguous instruction bytes to select the corresponding branch selector. For example, if a group of contiguous instruction bytes is 16 bytes, then decoder 84 decodes the four least significant bits of the fetch address. In this manner, a branch selector is chosen.

The selected branch selector is used to provide selection controls to branch prediction multiplexor 76, sequential/return multiplexor 78, and final prediction multiplexor 80. In one embodiment, the encoding of the branch selector can be used directly as the multiplexor select controls. In other embodiments, a logic block may be inserted between branch selector multiplexor 74 and multiplexors 76, 78, and 80. For the embodiment shown, branch selectors comprise two bits. One bit of the selected branch selector provides the selection control for prediction multiplexor 76 and sequential/return selector 78. The other bit provides a selection control for final prediction multiplexor 80. A branch prediction is thereby selected from the multiple branch predictions stored in branch prediction storage 70 corresponding to the group of contiguous instruction bytes being fetched, the sequential address of the group of contiguous instruction bytes sequential to the group of

contiguous instruction bytes being fetched, and a return stack address from a return stack structure. It is noted that multiplexors 76, 78, and 80 may be combined into a single 4 to 1 multiplexor for which the selected branch selector provides selection controls to select between the two branch predictions from branch prediction storage 70, the sequential address, and the return address.

5 The return stack structure (not shown) is used to store return addresses corresponding to subroutine call instructions previously fetched by microprocessor 10. In one embodiment, the branch predictions stored by branch prediction storage 70 include an indication that the branch prediction corresponds to a subroutine call instruction. Subroutine call instructions are a subset of branch instructions which save the address of the sequential instruction (the return address) in addition to redirecting the instruction stream to the target address  
10 of the subroutine call instruction. For example, in the x86 microprocessor architecture, the subroutine call instruction (CALL) pushes the return address onto the stack indicated by the ESP register.

A subroutine return instruction is another subset of the branch instructions. The subroutine return instruction uses the return address saved by the most recently executed subroutine call instruction as a target address. Therefore, when a branch prediction includes an indication that the branch prediction corresponds to  
15 a subroutine call instruction, the sequential address to the subroutine call instruction is placed at the top of the return stack. When a subroutine return instruction is encountered (as indicated by a particular branch selector encoding), the address nearest the top of the return stack which has not previously been used as a prediction is used as the prediction of the address. The address nearest the top of the return stack which has not previously been used as a prediction is conveyed by the return stack upon return address bus 92 (along with the predicted  
20 way of the return address, provided to the return stack similar to its provision upon way select bus 88. Branch prediction unit 14 informs the return stack when the return address is selected as the prediction. Additional details regarding an exemplary return stack structure may be found in the commonly assigned, co-pending patent application entitled: "Speculative Return Address Prediction Unit for a Superscalar Microprocessor", Serial No. 08/550,296, filed October 30, 1995 by Mahalingaiah, et al. The disclosure of the referenced patent  
25 application is incorporated herein by reference in its entirety.

The sequential address is provided by instruction cache 16. The sequential address identifies the next group of contiguous instruction bytes within main memory to the group of instruction bytes indicated by the fetch address upon fetch address bus 86. It is noted that, according to one embodiment, a way prediction is supplied for the sequential address when the sequential address is selected. The way prediction may be  
30 selected to be the same as the way selected for the fetch address. Alternatively, a way prediction for the sequential address may be stored within branch prediction storage 70.

As mentioned above, update logic block 82 is configured to update a prediction block upon detection of a branch misprediction or upon detection of a branch instruction while predecoding the corresponding group of contiguous instruction bytes in prefetch/predecode unit 12. The prediction block corresponding to  
35 each branch prediction is stored in update logic block 82 as the prediction is performed. A branch tag is conveyed along with the instructions being fetched (via a branch tag bus 102), such that if a misprediction is detected or a branch instruction is detected during predecoding, the corresponding prediction block can be identified via the branch tag. In one embodiment, the prediction block as shown in Fig. 5 is stored, as well as

the index of the fetch address which caused the prediction block to be fetched and the way in which the prediction block is stored.

When a branch misprediction is detected, the corresponding branch tag is provided upon mispredict bus 98 from either the functional unit 24 which executes the branch instruction or from decode units 20. If decode units 20 provide the branch tag, then the misprediction is of the previously undetected type (e.g. there are more branch instructions in the group than can be predicted using the corresponding branch predictions). Decode units 20 detect mispredictions of unconditional branch instructions (i.e. branch instructions which always select the target address). Functional units 24 may detect a misprediction due to a previously undetected conditional branch instruction or due to an incorrect taken/not-taken prediction. Update logic 82 selects the corresponding prediction block out of the aforementioned storage. In the case of a previously undetected branch instruction, one of the branch predictions within the prediction block is assigned to the previously undetected branch instruction. According to one embodiment, the algorithm for selecting one of the branch predictions to store the branch prediction for the previously undetected branch instruction is as follows: If the branch instruction is a subroutine return instruction, the branch selector for the instruction is selected to be the value indicating the return stack. Otherwise, a branch prediction which is currently predicted not-taken is selected. If each branch prediction is currently predicted-taken, then a branch prediction is randomly selected. The branch selector for the new prediction is set to indicate the selected branch prediction. Additionally, the branch selectors corresponding to bytes between the first branch instruction prior to the newly detected branch instruction and the newly detected branch instruction are set to the branch selector corresponding to the new prediction. Fig. 7 below describes one method for updating the branch selectors. For a mispredicted taken prediction which causes the prediction to become predicted not-taken, the branch selectors corresponding to the mispredicted prediction are set to the branch selector corresponding to the byte subsequent to the mispredicted branch instruction. In this manner, a prediction for a subsequent branch instruction will be used if the instructions are fetched again at a later clock cycle.

When prefetch/predecode unit 12 detects a branch instruction while predecoding a group of contiguous instruction bytes, prefetch/predecode unit 12 provides the branch tag for the group of contiguous instruction bytes if the predecoding is performed because invalid predecode information is stored in the instruction cache for the cache line (case (i)). Alternatively, if the predecoding is being performed upon a cache line being fetched from the main memory subsystem, prefetch/predecode unit 12 provides the address of the group of contiguous instruction bytes being predecoded, the offset of the end byte of the branch instruction within the group, and the way of the instruction cache selected to store the group (case (ii)). In case (i), the update is performed similar to the branch misprediction case above. In case (ii), there is not yet a valid prediction block stored in branch prediction storage 70 for the group of instructions. For this case, update logic block 82 initializes the branch selectors prior to the detected branch to the branch selector selected for the detected branch. Furthermore, the branch selectors subsequent to the detected branch are initialized to the sequential value. Alternatively, each of the branch selectors may be initialized to sequential when the corresponding cache line in instruction cache 16 is allocated, and subsequently updated via detection of a branch instructions during predecode in a manner similar to case (i).

Upon generation of an update, update logic block 82 conveys the updated prediction block, along with the fetch address index and corresponding way, upon update bus 96 for storage in branch prediction storage 70. It is noted that, in order to maintain branch prediction storage 70 as a single ported storage, branch prediction storage 70 may employ a branch holding register. The updated prediction information is stored into the branch holding register and updated into the branch prediction storage upon an idle cycle on fetch address bus 86. An exemplary cache holding register structure is described in the commonly assigned, co-pending patent application entitled: "Delayed Update Register for an Array", Serial No. 08/481,914, filed June 7, 1995, by Tran, et al., incorporated herein by reference in its entirety.

It is noted that a correctly predicted branch instruction may result in an update to the corresponding branch prediction as well. A counter indicative of previous executions of the branch instruction (used to form the taken/not-taken prediction of the branch instruction) may need to be incremented or decremented, for example. Such updates are performed upon retirement of the corresponding branch prediction. Retirement is indicated via a branch tag upon retire tag bus 104 from reorder buffer 32.

It is noted that the structure of Fig. 4 may be further accelerated through the use of a predicted branch selector. The predicted branch selector is stored with each prediction block and is set to the branch selector selected in a previous fetch of the corresponding group of contiguous instruction bytes. The predicted branch selector is used to select the branch prediction, removing branch selector multiplexor 74 from the path of branch prediction generation. Branch selector multiplexor 74 is still employed, however, to verify the selected branch selector is equal to the predicted branch selector. If the selected branch selector and the predicted branch selector are not equal, then the selected branch selector is used to provide the correct branch prediction during the succeeding clock cycle and the fetch of the incorrect branch prediction is cancelled.

Turning now to Fig. 5, an exemplary prediction block 110 employed by one embodiment of the branch prediction unit 14 as shown in Fig. 4 is shown. Prediction block 110 includes a set of branch selectors 112, a first branch prediction (BP1) 114, and a second branch prediction (BP2) 116. Set of branch selectors 112 includes a branch selector for each byte of the group of contiguous instruction bytes corresponding to prediction block 110.

First branch prediction 114 is shown in an exploded view in Fig. 5. Second branch prediction 116 is configured similarly. First branch prediction 114 includes an index 118 for the cache line containing the target address, and a way selection 120 for the cache line as well. According to one embodiment, index 118 includes the offset portion of the target address, as well as the index. Index 118 is concatenated with the tag of the way indicated by way selection 120 to form the branch prediction address. Additionally, a prediction counter 122 is stored for each branch prediction. The prediction counter is incremented each time the corresponding branch instruction is executed and is taken, and is decremented each time the corresponding branch instruction is executed and is not-taken. The most significant bit of the prediction counter is used as the taken/not-taken prediction. If the most significant bit is set, the branch instruction is predicted taken. Conversely, the branch instruction is predicted not-taken if the most significant bit is clear. In one embodiment, the prediction counter is a two bit saturating counter. The counter saturates when incremented at binary '11' and saturates when decremented at a binary '01'. In another embodiment, the prediction counter is a single bit which indicates a strong (a binary one) or a weak (a binary zero) taken prediction. If a strong taken prediction is mispredicted, it

becomes a weak taken prediction. If a weak taken prediction is mispredicted, the branch becomes predicted not taken and the branch selector is updated (i.e. the case of a mispredicted branch that becomes not-taken). Finally, a call bit 124 is included in first branch prediction 114. Call bit 124 is indicative, when set, that the corresponding branch instruction is a subroutine call instruction. If call bit 124 is set, the current fetch address and way are stored into the return stack structure mentioned above.

Turning next to Fig. 6, a table 130 illustrating an exemplary branch selector encoding is shown. A binary encoding is listed (most significant bit first), followed by the branch prediction which is selected when the branch selector is encoded with the corresponding value. As table 130 illustrates, the least significant bit of the branch selector can be used as a selection control for branch prediction multiplexor 76 and sequential/return multiplexor 78. If the least significant bit is clear, then the first branch prediction is selected by branch prediction multiplexor 76 and the sequential address is selected by sequential/return multiplexor 78. On the other hand, the second branch prediction is selected by branch prediction multiplexor 76 and the return address is selected by sequential/return multiplexor if the least significant bit is clear. Furthermore, the most significant bit of the branch selector can be used as a selection control for final prediction multiplexor 80. If the most significant bit is set, the output of branch prediction multiplexor 76 is selected. If the most significant bit is clear, the output of sequential/return multiplexor 78 is selected.

Turning now to Fig. 7, a flow chart depicting the steps employed to update the branch selectors of a group of contiguous instruction bytes in response to a mispredicted branch instruction is shown. Updating due to a branch instruction discovered during predecoding may be performed similarly. The misprediction may be the result of detecting a branch instruction for which prediction information is not stored in branch prediction storage 70, or may be the result of an incorrect taken/not-taken prediction which causes the corresponding prediction counter to indicate not-taken.

Upon detection of the misprediction, branch prediction unit 14 uses an "end pointer": the offset of the end byte of the mispredicted branch instruction within the corresponding group of contiguous instruction bytes. Additionally, the prediction block is selected for update using the branch tag received in response to the misprediction. Branch prediction unit 14 decodes the end pointer into an update mask (step 140). The update mask comprises a binary digit for each byte within the group of continuous instruction bytes. Digits corresponding to bytes prior to and including the end byte of the branch instruction within the cache line are set, and the remaining digits are clear.

Branch prediction unit 14 identifies the current branch selector. For mispredicted taken/not-taken predictions, the current branch selector is the branch selector corresponding to the mispredicted branch instruction. For misprediction due to an undetected branch, the current branch selector is the branch selector corresponding to the end byte of the undetected branch instruction. The current branch selector is XNOR'd with each of the branch selectors to create a branch mask (step 142). The branch mask includes binary digits which are set for each byte having a branch selector which matches the current branch selector and binary digits which are clear for each byte having a branch selector which does not match the current branch selector.

The update mask created in step 140 and the branch mask created in step 142 are subsequently ANDed, producing a final update mask (step 144). The final update mask includes binary digits which are set for each byte of the group of contiguous instruction bytes which is to be updated to the new branch selector.

For a mispredicted taken branch, the new branch selector is the branch selector of the byte subsequent to the end byte of the mispredicted taken branch instruction. For an undetected branch, the new branch selector is the branch selector indicating the branch prediction storage assigned to the previously undetected branch by update logic block 82.

5 An extended mask is also generated (steps 146 and 148). The extended mask indicates which branch selectors are to be erased because the branch prediction corresponding to the branch selector has been reallocated to the newly discovered branch instruction or because the branch prediction now indicates not taken. The extended mask is generated by first creating a second branch mask similar to the branch mask, except using the new branch selector instead of the current branch selector (i.e. the mask is created by  
10 XNORing the branch selectors corresponding to the cache line with the new branch selector (step 146)). The resulting mask is then ANDed with the inversion of the final update mask to create the extended mask (step 148). Branch selectors corresponding to bits in the extended mask which are set are updated to indicate the branch selector of the byte immediately subsequent to the last byte for which a bit in the extended mask is set. In this manner, the branch prediction formerly indicated by the branch selector is erased and replaced with the  
15 following branch selector within the cache line. During a step 150, the branch selectors are updated in response to the final update mask and the extended mask.

Turning now to Fig. 8, an example of the update of the branch selectors using the steps shown in the flowchart of Fig. 7 is shown. Each byte position is listed (reference number 160), followed by a set of branch selectors prior to update (reference number 162). In the initial set of branch selectors 162, a subroutine return  
20 instruction ends at byte position 1 as well as a first branch instruction ending at byte position 8 (indicated by branch selector number 3) and a second branch instruction ending at byte position 11 (indicated by branch selector number 2).

For the example of Fig. 8, a previously undetected branch instruction is detected ending at byte position 6. The second branch prediction is selected to represent the branch prediction for the previously  
25 undetected branch instruction. The update mask is generated as shown at reference number 164, given the end pointer of the previously undetected branch instruction is byte position 6. Since the example is a case of misprediction due to a previously undetected branch instruction and the branch selector at byte position 6 is "3", the current branch selector is "3". The XNORing of the current branch selector with the initial branch selectors 162 yields the branch mask depicted at reference number 166. The subsequent ANDing of the  
30 update mask and the branch mask yields the final update mask shown at reference number 168. As indicated by final update mask 168, byte positions 2 through 6 are updated to the new branch selector.

The second branch mask is generated by XNORing the new branch selector with initial branch selectors 162 (reference number 170). The new branch selector is "3", so second branch mask 170 is equal to  
35 branch mask 166 in this example. ANDing branch mask 170 with the logical inversion of the final update mask 168 produces the extended mask shown at reference number 172. As extended mask 172 indicates, byte positions 7 and 8 are to be updated to indicate the first branch prediction, since the second branch prediction has been assigned to the branch instruction ending at byte position 6 and the branch instruction represented by the first branch prediction ends at byte 11. An updated set of branch selectors is shown at reference number 174. The updated set of branch selectors at reference number 174 reflects choosing the branch prediction

corresponding to branch selector "3" for storing branch prediction information corresponding to the previously undetected branch instruction.

Turning next to Fig. 9, a second example of the update of the branch selectors using the steps shown in the flowchart of Fig. 7 is shown. Similar to the example of Fig. 8, each byte position is listed (reference number 160), followed by a set of branch selectors prior to update (reference number 162). In the initial set of branch selectors 162, a subroutine return instruction ends at byte position 1 as well as a first branch instruction ending at byte position 8 (indicated by branch selector number 3) and a second branch instruction ending at byte position 11 (indicated by branch selector number 2).

For the example of Fig. 9, a previously undetected branch instruction is again detected ending at byte position 6. However, the first branch prediction is selected to represent the branch prediction for the previously undetected branch instruction (as opposed to the second branch prediction as shown in Fig. 8). Since the misprediction is at the same byte position as Fig. 8, the same update mask, branch mask, and final update mask are generated as in Fig. 8 (reference numbers 164, 166, and 168).

The second branch mask is generated by XNORing the new branch selector with initial branch selectors 162 (reference number 180). The new branch selector is "2" in this example, so second branch mask 180 indicates byte positions 9 through 11. ANDing branch mask 180 with the logical inversion of the final update mask 168 produces the extended mask shown at reference number 182. As extended mask 182 indicates, byte positions 9 through 11 are to be updated to indicate the branch prediction following byte position 11 (i.e. the sequential branch prediction), since the first branch prediction has been assigned to the branch instruction ending at byte position 6 and the branch instruction represented by the second branch prediction ends at byte 8. An updated set of branch selectors is shown at reference number 184. The updated set of branch selectors at reference number 184 reflects choosing the branch prediction corresponding to branch selector "2" for storing branch prediction information corresponding to the previously undetected branch instruction.

Turning now to Fig. 10, a third example of the update of the branch selectors using the steps shown in the flowchart of Fig. 7 is shown. Similar to the example of Fig. 8, each byte position is listed (reference number 160), followed by a set of branch selectors prior to update (reference number 162). In the initial set of branch selectors 162, a subroutine return instruction ends at byte position 1 as well as a first branch instruction ending at byte position 8 (indicated by branch selector number 3) and a second branch instruction ending at byte position 11 (indicated by branch selector number 2).

For the example of Fig. 10, a the branch instruction ending at byte position 8 is mispredicted as taken, and the ensuing update of the second branch prediction causes the prediction counter to indicate not-taken. Since the branch prediction is not taken, the branch selectors indicating the branch prediction should be updated to indicate the subsequent branch instruction (or update to indicate sequential, if there is no subsequent branch instruction within the group of contiguous instruction bytes). In cases in which a branch prediction becomes not-taken, the end pointer of the "new" branch instruction is invalid, since there is no newly detected branch instruction. Therefore, the update mask is generated as all zero (reference number 190). Since the current branch selector is "3", the branch mask is generated as shown at reference number 191. Therefore, the final update mask (reference number 192) is all zeros.

The second branch mask is generated by XNORing the new branch selector with initial branch selectors 162 (reference number 194). The new branch selector is set to "3" in this example, such that each of the branch selectors coded to "3" are indicated by second branch mask 194. ANDing branch mask 180 with the logical inversion of the final update mask 192 produces the extended mask shown at reference number 196. As extended mask 196 indicates, byte positions 2 through 8 are to be updated to indicate the branch prediction following byte position 8 (i.e. the first branch prediction), since the first branch prediction is assigned to the branch instruction ending at byte position 11. An updated set of branch selectors is shown at reference number 198. The updated set of branch selectors at reference number 198 reflects deleting branch selector "3" from the set of branch selectors corresponding to the group of contiguous instruction bytes, since the first branch prediction is not storing a predicted-taken branch prediction.

As Fig. 10 illustrates, the procedure for removing a branch selector when a prediction indicates not-taken is similar to the procedure for reassigning a branch prediction. The differences in the two procedures are that the update mask for removing a branch selector is always generated as zeros, and the current branch selector is provided as the "new" branch selector in order to generate the extended mask.

It is noted that, although the preceding discussion has focused on an embodiment in which a variable byte-length instruction set is employed (e.g. the x86 instruction set), branch selectors may be employed in branch prediction mechanisms for fixed byte length instruction sets as well. An embodiment for fixed byte length instruction sets may store a branch selector for each instruction, since the instructions are stored at constant offsets within cache lines or groups of contiguous instruction bytes.

It is further noted that, although the embodiment above shows multiple branch predictions per group of contiguous instruction bytes, branch selectors may be employed even when only one branch prediction is stored for each group. The branch selectors in this case may be a single bit. If the bit is set, then the branch prediction is selected. If the bit is clear, then the sequential prediction is selected.

It is noted that, as referred to above, a previously undetected branch instruction is a branch instruction represented by none of the branch predictions within the corresponding prediction block. The previously undetected branch instruction may be previously undetected (i.e. not executed since the corresponding cache line was stored into instruction cache 16). Alternatively, the branch prediction corresponding to the previously undetected branch instruction may have been reassigned to a different branch instruction within the corresponding group of contiguous instruction bytes.

Turning now to Fig. 11, a computer system 200 including microprocessor 10 is shown. Computer system 200 further includes a bus bridge 202, a main memory 204, and a plurality of input/output (I/O) devices 206A-206N. Plurality of I/O devices 206A-206N will be collectively referred to as I/O devices 206. Microprocessor 10, bus bridge 202, and main memory 204 are coupled to a system bus 208. I/O devices 206 are coupled to an I/O bus 210 for communication with bus bridge 202.

Bus bridge 202 is provided to assist in communications between I/O devices 206 and devices coupled to system bus 208. I/O devices 206 typically require longer bus clock cycles than microprocessor 10 and other devices coupled to system bus 208. Therefore, bus bridge 202 provides a buffer between system bus 208 and input/output bus 210. Additionally, bus bridge 202 translates transactions from one bus protocol to another. In one embodiment, input/output bus 210 is an Enhanced Industry Standard Architecture (EISA) bus

and bus bridge 202 translates from the system bus protocol to the EISA bus protocol. In another embodiment, input/output bus 210 is a Peripheral Component Interconnect (PCI) bus and bus bridge 202 translates from the system bus protocol to the PCI bus protocol. It is noted that many variations of system bus protocols exist.

Microprocessor 10 may employ any suitable system bus protocol.

5 I/O devices 206 provide an interface between computer system 200 and other devices external to the computer system. Exemplary I/O devices include a modem, a serial or parallel port, a sound card, etc. I/O devices 206 may also be referred to as peripheral devices. Main memory 204 stores data and instructions for use by microprocessor 10. In one embodiment, main memory 204 includes at least one Dynamic Random Access Memory (DRAM) and a DRAM memory controller.

10 It is noted that although computer system 200 as shown in Fig. 11 includes one bus bridge 202, other embodiments of computer system 200 may include multiple bus bridges 202 for translating to multiple dissimilar or similar I/O bus protocols. Still further, a cache memory for enhancing the performance of computer system 200 by storing instructions and data referenced by microprocessor 10 in a faster memory storage may be included. The cache memory may be inserted between microprocessor 10 and system bus 208, 15 or may reside on system bus 208 in a "lookaside" configuration.

Although various components above have been described as multiplexors, it is noted that multiple multiplexors, in series or in parallel, may be employed to perform the selection represented by the multiplexors shown.

It is still further noted that the present discussion may refer to the assertion of various signals. As 20 used herein, a signal is "asserted" if it conveys a value indicative of a particular condition. Conversely, a signal is "deasserted" if it conveys a value indicative of a lack of a particular condition. A signal may be defined to be asserted when it conveys a logical zero value or, conversely, when it conveys a logical one value. Additionally, various values have been described as being discarded in the above discussion. A value may be discarded in a number of manners, but generally involves modifying the value such that it is ignored by logic circuitry which 25 receives the value. For example, if the value comprises a bit, the logic state of the value may be inverted to discard the value. If the value is an n-bit value, one of the n-bit encodings may indicate that the value is invalid. Setting the value to the invalid encoding causes the value to be discarded. Additionally, an n-bit value may include a valid bit indicative, when set, that the n-bit value is valid. Resetting the valid bit may comprise discarding the value. Other methods of discarding a value may be used as well.

30 Table 1 below indicates fast path, double dispatch, and MROM instructions for one embodiment of microprocessor 10 employing the x86 instruction set:

**Table 1: x86 Fast Path, Double Dispatch, and MROM Instructions**

<u>X86 Instruction</u>	<u>Instruction Category</u>
AAA	MROM
AAD	MROM
AAM	MROM
AAS	MROM
ADC	fast path
ADD	fast path
AND	fast path
ARPL	MROM
BOUND	MROM
BSF	fast path

BSR	fast path
BSWAP	MROM
BT	fast path
BTC	fast path
BTR	fast path
BTS	fast path
CALL	fast path/double dispatch
CBW	fast path
CWDE	fast path
CLC	fast path
CLD	fast path
CLI	MROM
CLTS	MROM
CMC	fast path
CMP	fast path
CMPS	MROM
CMPSB	MROM
CMPSW	MROM
CMPSD	MROM
CMPXCHG	MROM
CMPXCHG8B	MROM
CPUID	MROM
CWD	MROM
CWQ	MROM
DDA	MROM
DAS	MROM
DEC	fast path
DIV	MROM
ENTER	MROM
HLT	MROM
IDIV	MROM
IMUL	double dispatch
IN	MROM
INC	fast path
INS	MROM
INSB	MROM
INSW	MROM
INSD	MROM
INT	MROM
INTO	MROM
INVD	MROM
INVLPG	MROM
IRET	MROM
IRETD	MROM
Jcc	fast path
JCXZ	double dispatch
JECXZ	double dispatch
JMP	fast path
LAHF	fast path
LAR	MROM
LDS	MROM
LES	MROM
LFS	MROM
LGS	MROM
LSS	MROM
LEA	fast path
LEAVE	double dispatch
LGDT	MROM
LIDT	MROM

LLDT	MROM
LMSW	MROM
LODS	MROM
LODSB	MROM
LODSW	MROM
LOSD	MROM
LOOP	double dispatch
LOOPcond	MROM
LSL	MROM
LTR	MROM
MOV	fast path
MOVCC	fast path
MOV.CR	MROM
MOV.DR	MROM
MOVS	MROM
MOVSB	MROM
MOVSW	MROM
MOVSD	MROM
MOVSX	fast path
MOVZX	fast path
MUL	double dispatch
NEG	fast path
NOP	fast path
NOT	fast path
OR	fast path
OUT	MROM
OUTS	MROM
OUTSB	MROM
OUTSW	MROM
OUTSD	MROM
POP	double dispatch
POPA	MROM
POPAD	MROM
POPF	MROM
POPFD	MROM
PUSH	fast path/double dispatch
PUSHA	MROM
PUSHAD	MROM
PUSHF	fast path
PUSHFD	fast path
RCL	MROM
RCR	MROM
ROL	fast path
ROR	fast path
RDMSR	MROM
REP	MROM
REPE	MROM
REPZ	MROM
REPNE	MROM
REPNZ	MROM
RET	double dispatch
RSM	MROM
SAHF	fast path
SAL	fast path
SAR	fast path
SHL	fast path
SHR	fast path
SBB	fast path
SCAS	double dispatch

SCASB	MROM
SCASW	MROM
SCASD	MROM
SETcc	fast path
SGDT	MROM
SIDT	MROM
SHLD	MROM
SHRD	MROM
SLDT	MROM
SMSW	MROM
STC	fast path
STD	fast path
STI	MROM
STOS	MROM
STOSB	MROM
STOSW	MROM
STOSD	MROM
STR	MROM
SUB	fast path
TEST	fast path
VERR	MROM
VERW	MROM
WBINVD	MROM
WRMSR	MROM
XADD	MROM
XCHG	MROM
XLAT	fast path
XLATB	fast path
XOR	fast path

Note: Instructions including an SIB byte are also considered double dispatch instructions.

It is noted that a superscalar microprocessor in accordance with the foregoing may further employ the  
 5 latching structures as disclosed within the co-pending, commonly assigned patent application entitled  
 "Conditional Latching Mechanism and Pipelined Microprocessor Employing the Same", Serial No.  
 08/400,608 filed March 8, 1995, by Pflum et al. The disclosure of this patent application is incorporated  
 herein by reference in its entirety.

It is further noted that aspects regarding array circuitry may be found in the co-pending, commonly  
 10 assigned patent application entitled "High Performance Ram Array Circuit Employing Self-Time Clock  
 Generator for Enabling Array Access", Serial No. 08/473,103 filed June 7, 1995 by Tran. The disclosure of  
 this patent application is incorporated herein by reference in its entirety.

It is additionally noted that other aspects regarding superscalar microprocessors may be found in the  
 following co-pending, commonly assigned patent applications: "Linearly Addressable Microprocessor Cache",  
 15 Serial No. 08/146,381, filed Oct 29, 1993 by Witt; "Superscalar Microprocessor Including a High  
 Performance Instruction Alignment Unit", Serial No. 08/377,843, filed January 25, 1995 by Witt, et al; "A  
 Way Prediction Structure", Serial No. 08/522,181, filed August 31, 1995 by Roberts, et al; "A Data Cache  
 Capable of Performing Store Accesses in a Single Clock Cycle", Serial No. 08/521,627, filed August 31, 1995  
 by Witt, et al; "A Parallel and Scalable Instruction Scanning Unit", Serial No. 08/475,400, filed June 7, 1995  
 20 by Narayan; and "An Apparatus and Method for Aligning Variable-Byte Length Instructions to a Plurality of

Issue Positions", Serial No. 08/582,473, filed January 2, 1996 by Tran, et al. The disclosure of these patent applications are incorporated herein by reference in their entirety.

5 In accordance with the above disclosure, a branch prediction mechanism using branch selectors is described. The branch prediction mechanism quickly locates the branch prediction corresponding to a given fetch address by selecting the branch selector corresponding to the byte indicated by the given fetch address and selecting the branch prediction indicated by that branch selector. The branch prediction mechanism may be capable of operating at a higher frequency than previous branch prediction mechanisms.

10 Numerous variations and modifications will become apparent to those skilled in the art once the above disclosure is fully appreciated. It is intended that the following claims be interpreted to embrace all such variations and modifications.

**WHAT IS CLAIMED IS:**

1. A method for selecting a branch prediction corresponding to a group of contiguous instruction bytes including a plurality of instructions, comprising:

5

storing a plurality of branch selectors in a branch prediction storage, wherein at least one of said plurality of branch selectors corresponds to a first one of said plurality of instructions;

10

fetching said plurality of branch selectors using a fetch address identifying said group of contiguous instruction bytes;

selecting one of said plurality of branch selectors in response to said fetch address; and

15

selecting said branch prediction in response to said one of said plurality of said branch selectors.

2. The method as recited in claim 1 wherein said one of said plurality of branch selectors identifies a particular branch prediction to be selected if said first one of said plurality of instructions is fetched.

20

3. The method as recited in claim 2 further comprising concurrently fetching said group of contiguous instruction bytes with said plurality of branch selectors.

25

4. The method as recited in claim 3 wherein said particular branch prediction comprises one of a plurality of branch predictions stored in said branch prediction storage with respect to said group of contiguous instruction bytes.

5. The method as recited in claim 4 further comprising updating said plurality of branch predictions in response to a branch misprediction.

30

6. The method as recited in claim 5 wherein said updating comprises assigning one of said plurality of branch predictions stored in said branch prediction storage to a previously undetected branch instruction within said group of contiguous instruction bytes if said branch misprediction is a result of said previously undetected branch instruction.

35

7. The method as recited in claim 6 further comprising updating said plurality of branch selectors to reflect said previously undetected branch instruction and to delete branch selectors corresponding to a branch instruction previously predicted by said one of said plurality of branch predictions.

8. The method as recited in claim 3 wherein said particular branch prediction comprises a return stack address corresponding to a previously fetched subroutine call instruction.

9. The method as recited in claim 3 wherein said particular branch prediction comprises a sequential address corresponding to a second group of contiguous instruction bytes immediately subsequent to said group of contiguous instruction bytes within a main memory.

5

10. The method as recited in claim 3 wherein said selecting one of said plurality of branch selectors comprises decoding a plurality of least significant bits of said fetch address.

11. The method as recited in claim 10 wherein said plurality of least significant bits of said fetch address comprise an offset into said group of contiguous instruction bytes whereby any one of said group of contiguous instruction bytes is identifiable via a different value of said offset.

10

12. The method as recited in claim 3 wherein said plurality of instructions belong to a variable byte length instruction set.

15

13. The method as recited in claim 12 wherein each of said plurality of branch selectors corresponds to a different one of said group of contiguous instruction bytes.

14. A branch prediction apparatus, comprising:

20

a branch prediction storage coupled to receive a fetch address corresponding to a group of contiguous instruction bytes being fetched from an instruction cache, said branch prediction storage configured to store a plurality of branch selectors wherein at least one of said plurality of branch selectors corresponds to a first instruction within said group of contiguous instruction bytes; and

25

a selection mechanism coupled to said branch prediction storage to receives said plurality of branch selectors, said selection mechanism configured to select a particular one of said plurality of branch selectors in response to said fetch address.

30

15. The branch prediction apparatus as recited in claim 14 wherein said one of said plurality of branch selectors identifies a particular branch prediction to be selected if said first instruction is fetched.

16. The branch prediction apparatus as recited in claim 15 wherein said selection mechanism is configured to select said particular one of said plurality of branch selectors responsive to a plurality of least significant bits of said fetch address used to fetch said group of contiguous instruction bytes.

35

17. The branch prediction apparatus as recited in claim 16 wherein said group of contiguous instruction bytes comprises a portion of a cache line of an instruction cache storing said group of contiguous instruction bytes.

18. The branch prediction apparatus as recited in claim 17 wherein said portion of said cache line comprises one half of said cache line.

5 19. The branch prediction apparatus as recited in claim 16 wherein said particular branch prediction comprises one of a plurality of branch predictions stored in said branch prediction storage with respect to said group of contiguous instruction bytes.

10 20. The branch prediction apparatus as recited in claim 19 wherein at least one of said plurality of branch selectors identifies a sequential address as said particular branch prediction, wherein said sequential address corresponds to a second group of contiguous instruction bytes immediately subsequent to said group of contiguous instruction bytes within a main memory.

15 21. The branch prediction apparatus as recited in claim 20 wherein said one of said plurality of branch selectors identifies a return stack address as said particular branch prediction if a corresponding branch instruction within said group of contiguous instruction bytes is a subroutine return instruction.

20 22. The branch prediction apparatus as recited in claim 21 further comprising a second selection mechanism coupled to receive said particular one of said plurality of branch selectors and coupled to receive said plurality of branch predictions, said sequential address, and said return stack address, and wherein said second selection mechanism is configured to select an output branch prediction of said branch prediction apparatus in response to said particular one of said plurality of branch selectors.

25 23. The branch prediction apparatus as recited in claim 22 wherein said second selection mechanism is configured to use a first bit of said one of said branch selectors to select between one of said plurality of branch predictions and another one of said plurality of branch predictions and further selects between said sequential address and said return stack address.

30 24. The branch prediction apparatus as recited in claim 23 wherein said second selection mechanism is configured to use a second bit of said one of said branch selectors to select between addresses selected using said first bit.

25. A microprocessor comprising:

35 an instruction cache configured to store a plurality of cache lines of instruction bytes and to provide a group of instruction bytes upon receipt of a fetch address to an instruction processing pipeline of said microprocessor; and

a branch prediction unit coupled to said instruction cache and coupled to receive said fetch address, wherein said branch prediction unit is configured to store a plurality of branch selectors with respect to said group of instruction bytes and is configured to select one of said plurality of branch selectors in response to said fetch address.

5

26. The microprocessor as recited in claim 25 wherein said branch prediction unit is coupled to receive said fetch address concurrent with said instruction cache receiving said fetch address.

27. The microprocessor as recited in claim 26 wherein said one of said plurality of branch selectors identifies a branch prediction which is used as a subsequent fetch address by said instruction cache.

10

28. The microprocessor as recited in claim 27 wherein said one of said plurality of branch selectors is selected by decoding a plurality of least significant bits of said fetch address.

29. A branch prediction apparatus configured to select a branch selector, said branch selector corresponding to a first instruction within a group of contiguous instruction bytes.

15

30. A method for performing branch prediction comprising selecting a branch selector corresponding to a first instruction within a group of contiguous instruction bytes.

20

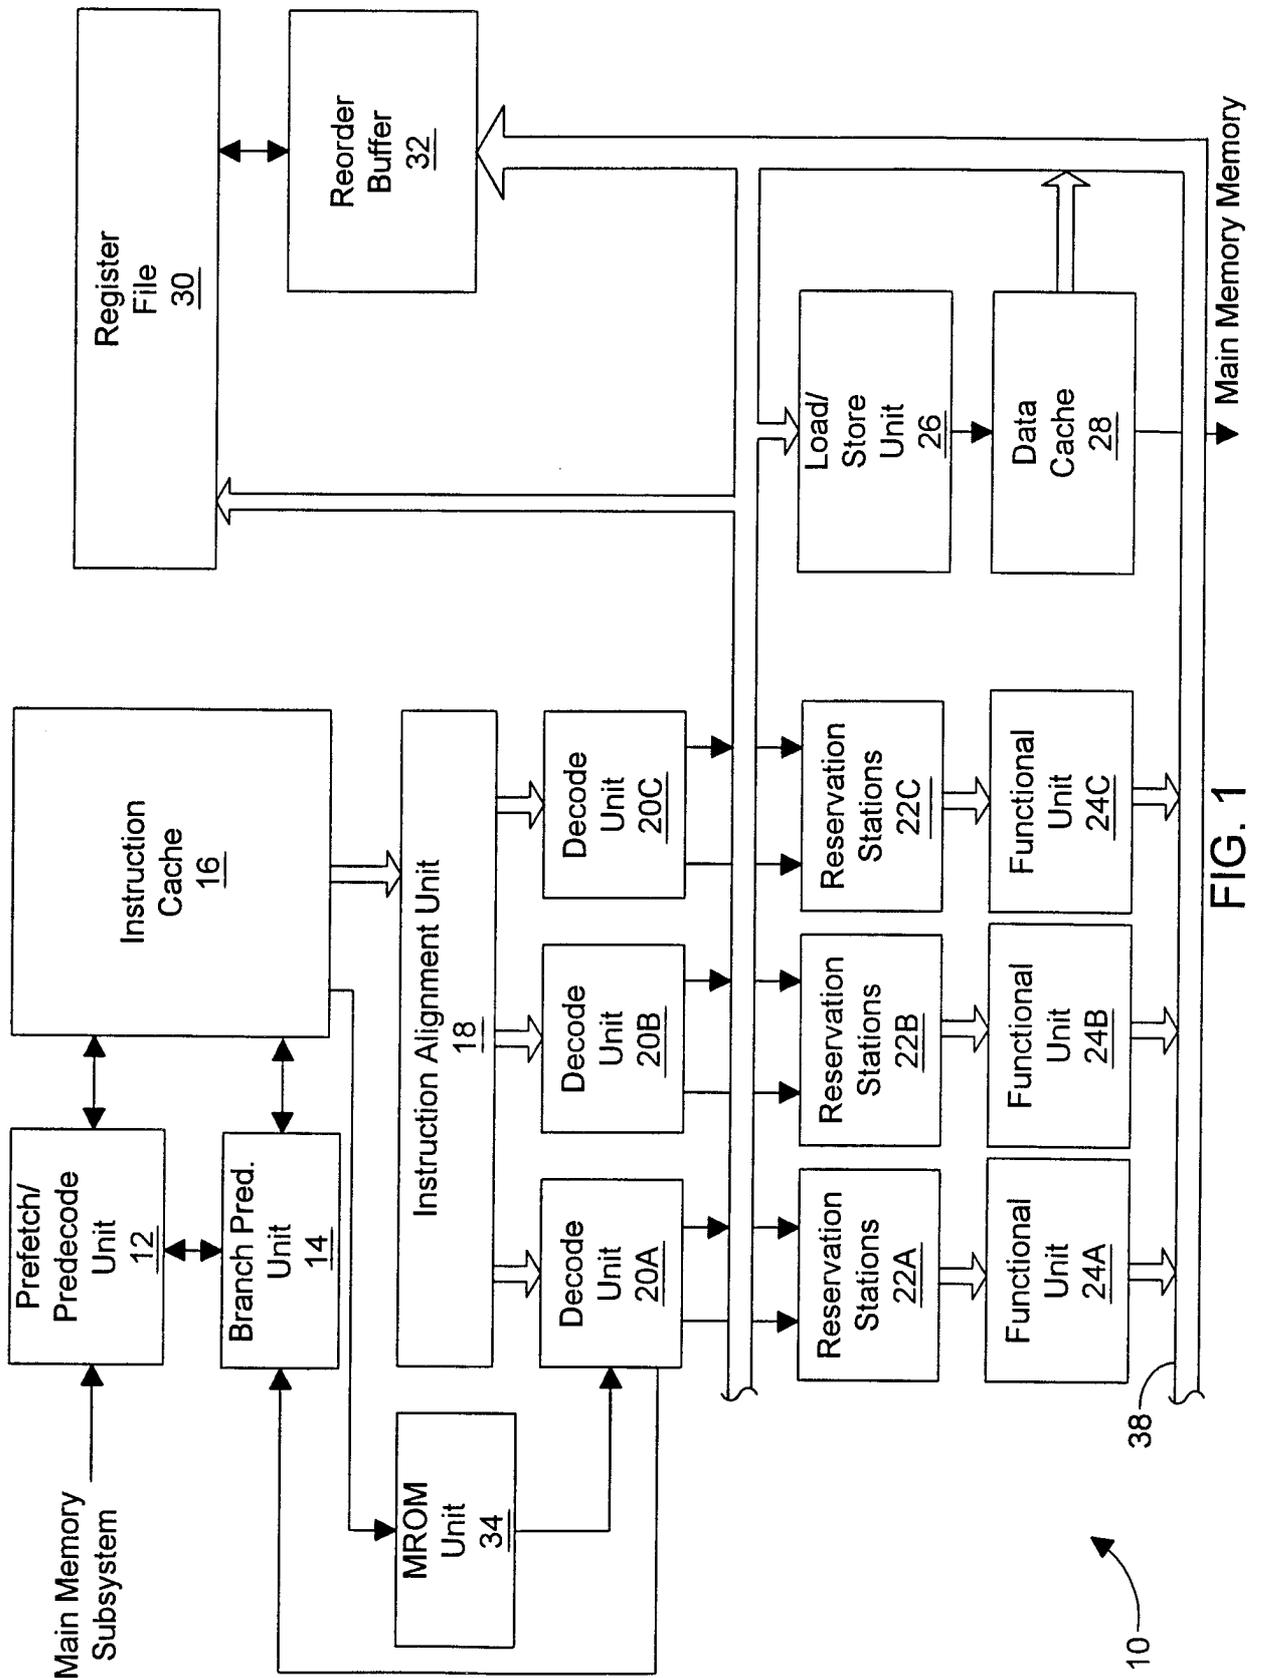


FIG. 1

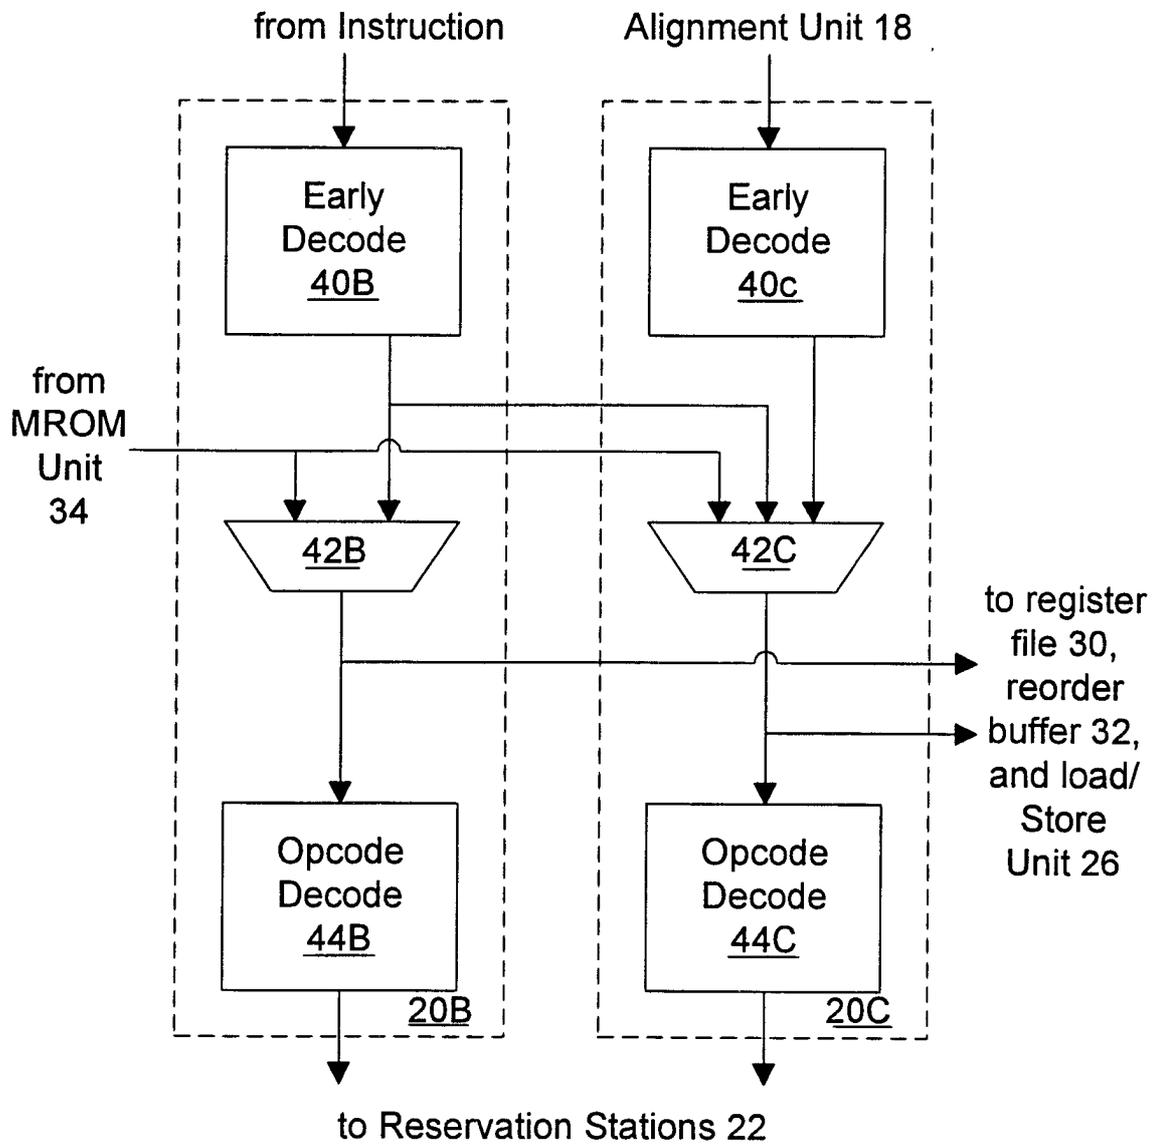


FIG. 2

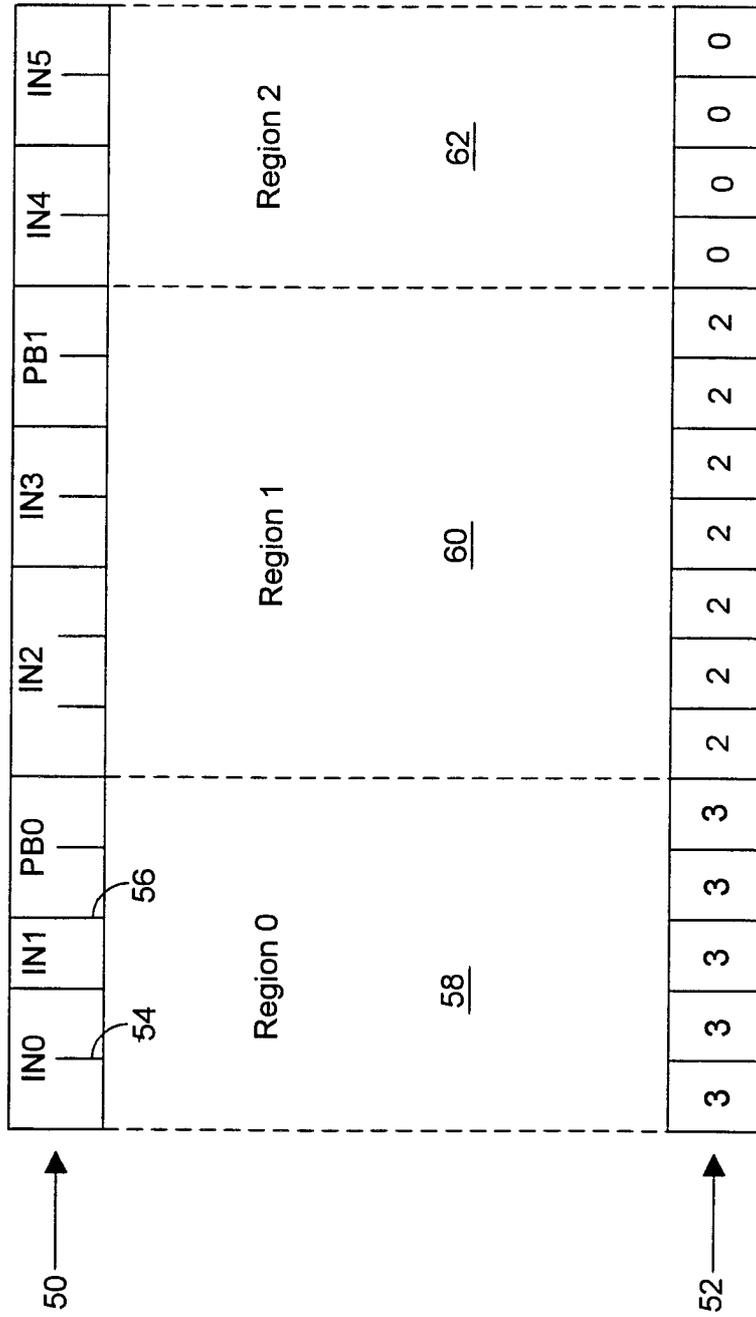


FIG. 3

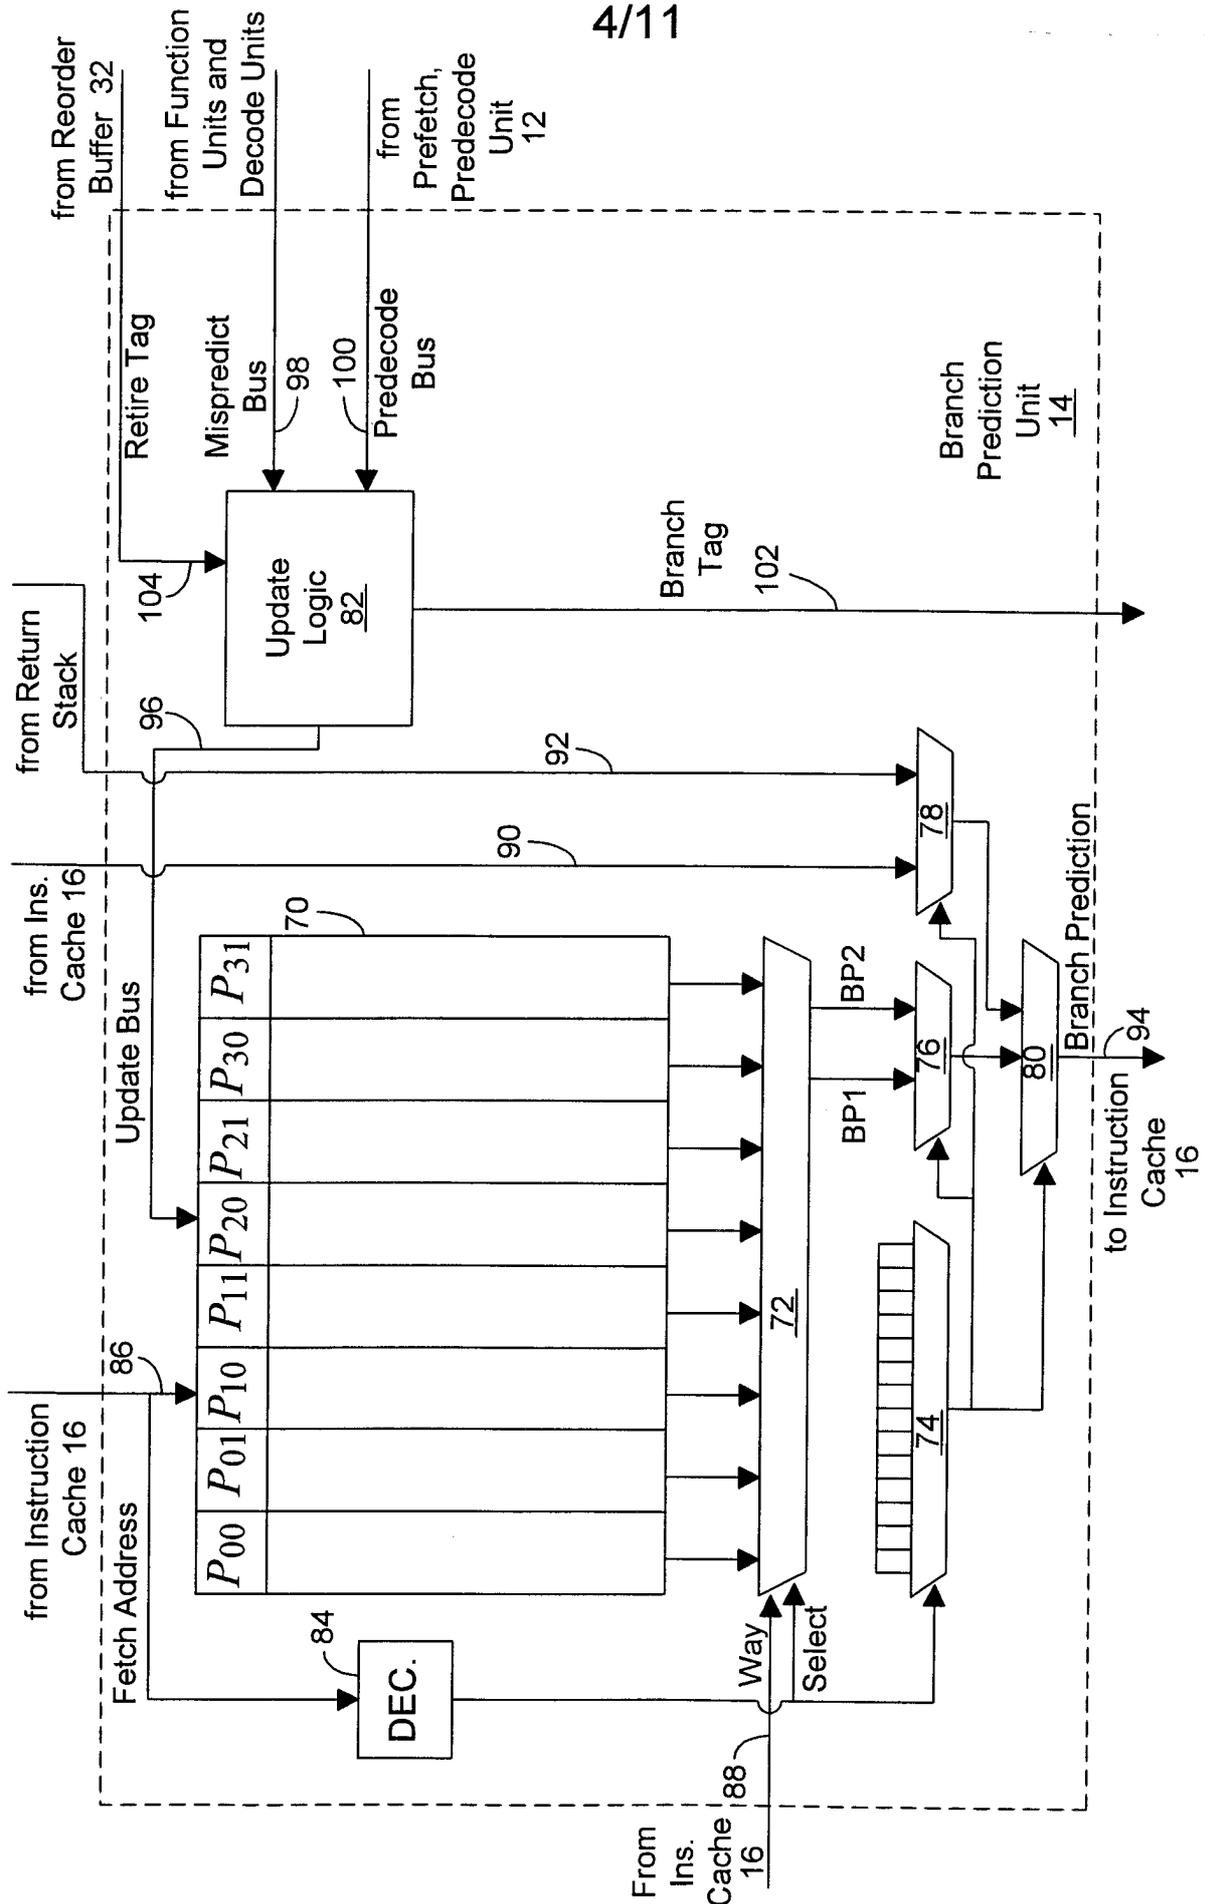


FIG. 4

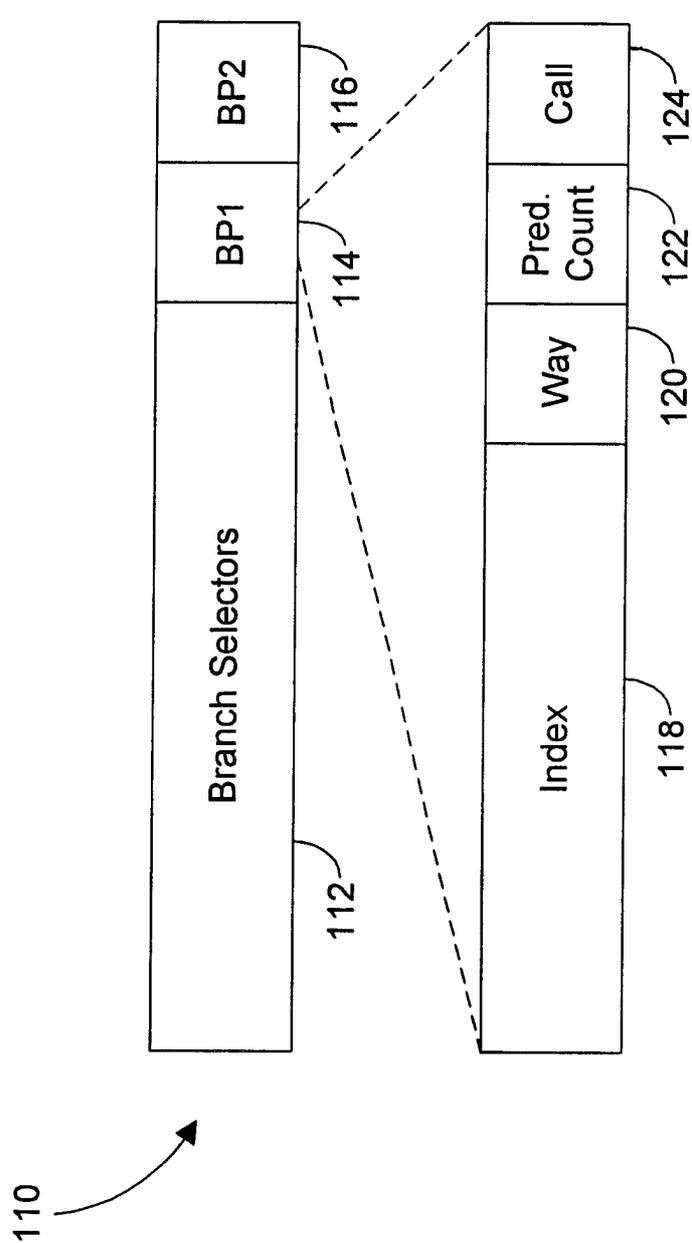
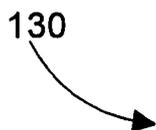


FIG. 5

130  


<u>Encoding (Binary)</u>	<u>Branch Prediction</u>
00	Sequential
01	Return Stack
10	Branch Prediction 1
11	Branch Prediction 2

FIG. 6

7/11

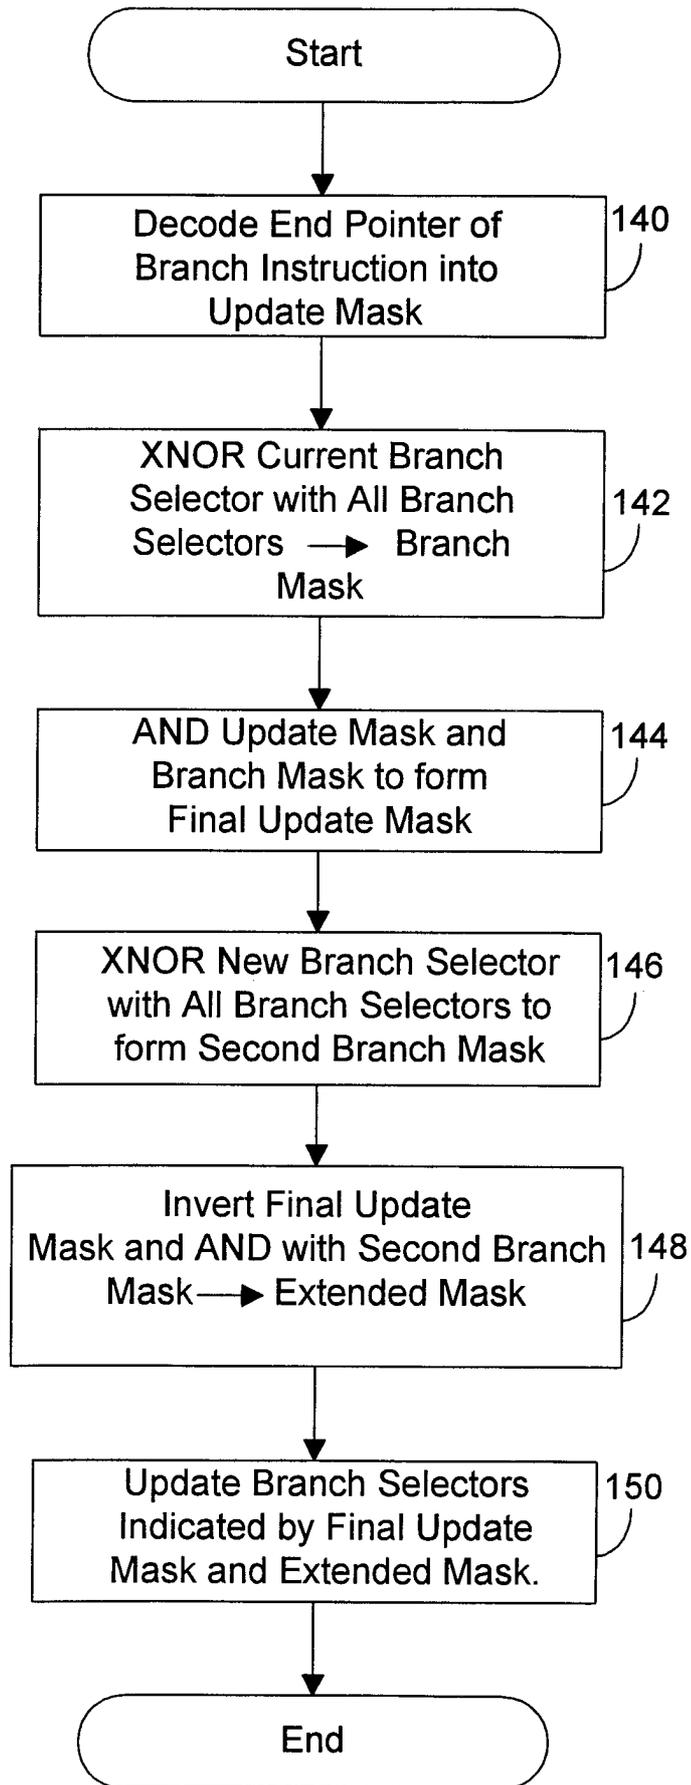


FIG. 7

Byte Position	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	160
Branch Selectors	1	1	3	3	3	3	3	3	3	2	2	2	0	0	0	0	162
Update Mask	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	164
Branch Mask	0	0	1	1	1	1	1	1	1	0	0	0	0	0	0	0	166
Final Update Mask	0	0	1	1	1	1	1	1	0	0	0	0	0	0	0	0	168
Second Branch Mask	0	0	1	1	1	1	1	1	1	0	0	0	0	0	0	0	170
Extended Mask	0	0	0	0	0	0	0	1	1	0	0	0	0	0	0	0	172
Updated Branch Selectors	1	1	3	3	3	3	3	2	2	2	2	2	0	0	0	0	174

Mispredicted Branch = Byte Position 6

New Branch Selector = 3

FIG. 8

Byte Position	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	160
Branch Selectors	1	1	3	3	3	3	3	3	3	2	2	2	0	0	0	0	162
Update Mask	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	164
Branch Mask	0	0	1	1	1	1	1	1	1	0	0	0	0	0	0	0	166
Final Update Mask	0	0	1	1	1	1	1	1	0	0	0	0	0	0	0	0	168
Second Branch Mask	0	0	0	0	0	0	0	0	0	1	1	1	0	0	0	0	180
Extended Mask	0	0	0	0	0	0	0	0	0	1	1	1	0	0	0	0	182
Updated Branch Selectors	1	1	2	2	2	2	2	2	3	0	0	0	0	0	0	0	184

Mispredicted Branch = Byte Position 6

New Branch Selector = 2

FIG. 9

Byte Position	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	160
Branch Selectors	1	1	3	3	3	3	3	3	3	2	2	2	0	0	0	0	162
Update Mask	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	190
Branch Mask	0	0	1	1	1	1	1	1	1	0	0	0	0	0	0	0	191
Final Update Mask	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	192
Second Branch Mask	0	0	1	1	1	1	1	1	1	0	0	0	0	0	0	0	194
Extended Mask	0	0	1	1	1	1	1	1	1	0	0	0	0	0	0	0	196
Updated Branch Selectors	1	1	2	2	2	2	2	2	2	2	2	2	0	0	0	0	198

Branch changes from Predicted  
taken to Predicted not taken

Branch Selector = 3

FIG. 10

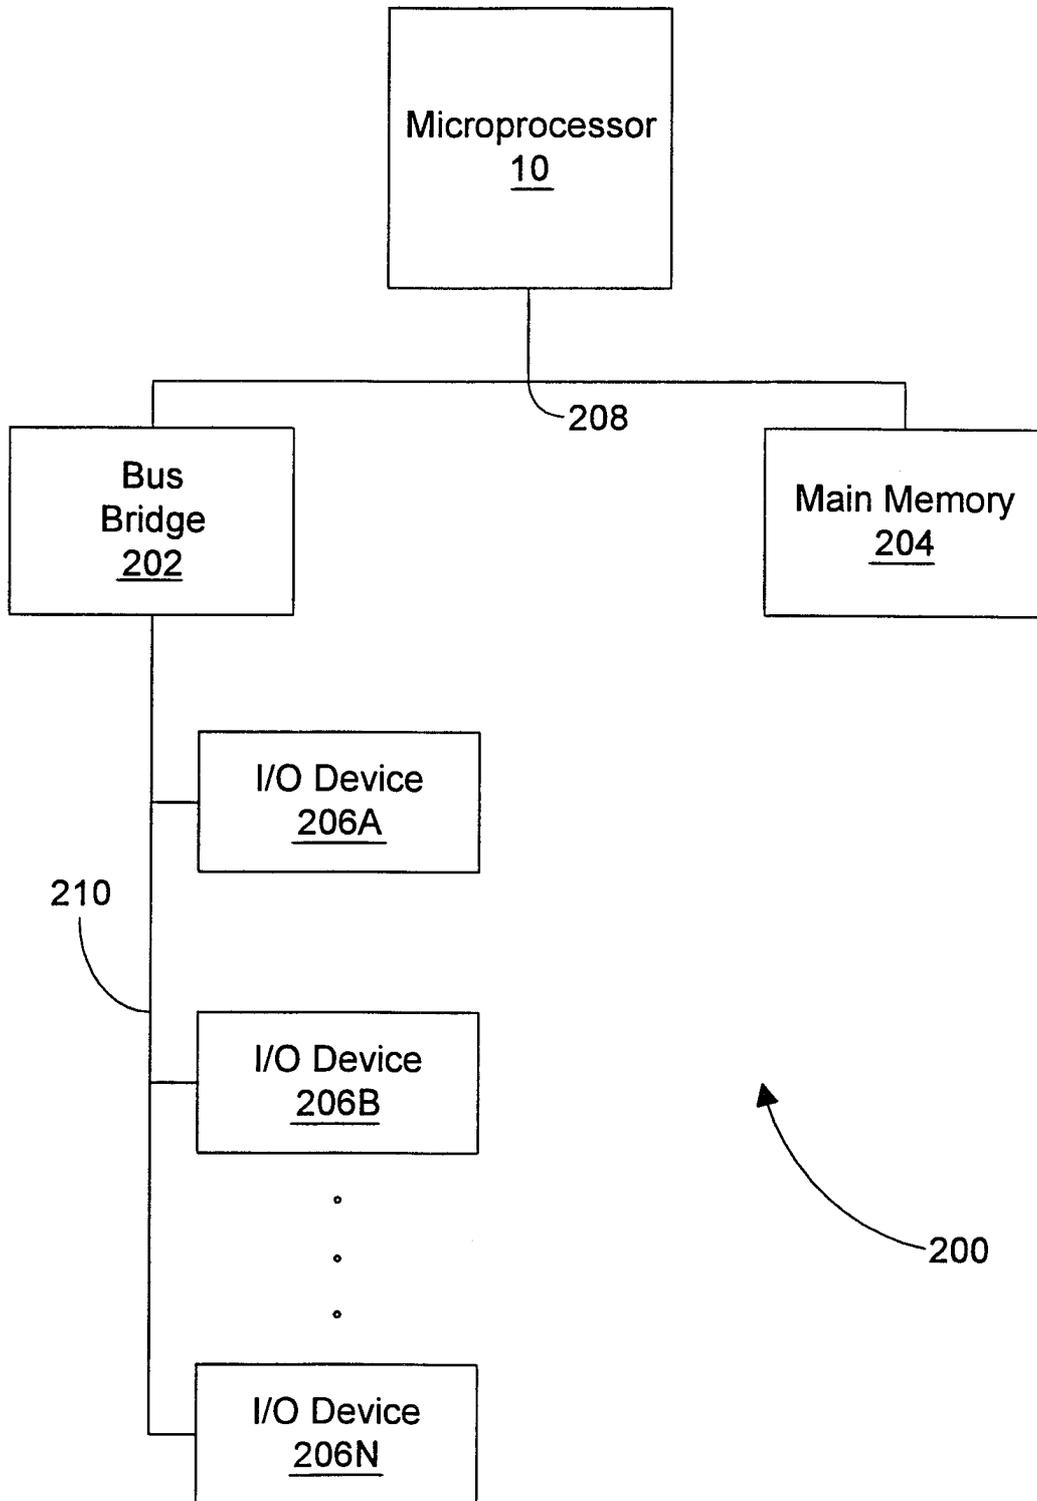


FIG. 11

# INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 97/21048

**A. CLASSIFICATION OF SUBJECT MATTER**

IPC 6 G06F9/38

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

**B. FIELDS SEARCHED**

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 G06F

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

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

**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	EP 0 605 876 A (IBM) 13 July 1994 see the whole document ---	1-30
A	EP 0 586 057 A (SUN MICROSYSTEMS INC) 9 March 1994 see page 4, line 45 - page 6, line 9 ---	1, 14, 25
A	GB 2 285 526 A (INTEL CORP) 12 July 1995 see page 21, line 5 - page 24, line 10 ---	4, 19
A	EP 0 199 947 A (IBM) 5 November 1986 see page 19 - page 21 ---	1, 14, 25
A	CALDER B ET AL: "NEXT CACHE LINE AND SET PREDICTION" COMPUTER ARCHITECTURE NEWS, vol. 23, no. 2, 1 May 1995, pages 287-296, XP000525181 -----	

Further documents are listed in the continuation of box C.

Patent family members are listed in annex.

° Special categories of cited documents :

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

20 March 1998

Date of mailing of the international search report

30/03/1998

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2  
NL - 2280 HV Rijswijk  
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,  
Fax: (+31-70) 340-3016

Authorized officer

Daskalakis, T

# INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 97/21048

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 0605876 A	13-07-94	US 5367703 A JP 2531495 B JP 6236270 A	22-11-94 04-09-96 23-08-94
EP 0586057 A	09-03-94	JP 6208463 A	26-07-94
GB 2285526 A	12-07-95	US 5574871 A DE 4447238 A JP 7210383 A US 5706492 A	12-11-96 13-07-95 11-08-95 06-01-98
EP 0199947 A	05-11-86	US 4679141 A CA 1244554 A DE 3682700 A JP 1656288 C JP 3018211 B JP 61250738 A	07-07-87 08-11-88 16-01-92 13-04-92 12-03-91 07-11-86