




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 9/28	A1	(11) International Publication Number: WO 95/18408 (43) International Publication Date: 6 July 1995 (06.07.95)
(21) International Application Number: PCT/US94/14929 (22) International Filing Date: 28 December 1994 (28.12.94) (30) Priority Data: 08/174,856 29 December 1993 (29.12.93) US 08/174,858 29 December 1993 (29.12.93) US (71) Applicant: UNISYS CORPORATION [US/US]; Township Line and Union Meeting Roads, P.O. Box 500 - C1SW19, Blue Bell, PA 19424 (US). (72) Inventor: FLORA, Laurence, P.; 28331 Cepin Drive, Valley Center, CA 92082 (US). (74) Agent: STARR, Mark, T.; Unisys Corporation, Township Line and Union Meeting Roads, P.O. Box 500 - C1SW19, Blue Bell, PA 19424 (US).		(81) Designated States: JP, KR, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published <i>With international search report.</i> <i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>

(54) Title: PIPELINED MICROINSTRUCTION APPARATUS AND METHODS WITH BRANCH PREDICTION AND SPECULATIVE STATE CHANGING

(57) Abstract

High speed instruction execution apparatus is disclosed which provides multistage pipelining and branch prediction in a manner which permits speculative changes of state to be made during execution of a predicted instruction before the correctness of the prediction has been determined.



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.

AT	Austria	GB	United Kingdom	MR	Mauritania
AU	Australia	GE	Georgia	MW	Malawi
BB	Barbados	GN	Guinea	NE	Niger
BE	Belgium	GR	Greece	NL	Netherlands
BF	Burkina Faso	HU	Hungary	NO	Norway
BG	Bulgaria	IE	Ireland	NZ	New Zealand
BJ	Benin	IT	Italy	PL	Poland
BR	Brazil	JP	Japan	PT	Portugal
BY	Belarus	KE	Kenya	RO	Romania
CA	Canada	KG	Kyrgyzstan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SI	Slovenia
CI	Côte d'Ivoire	LI	Liechtenstein	SK	Slovakia
CM	Cameroon	LK	Sri Lanka	SN	Senegal
CN	China	LU	Luxembourg	TD	Chad
CS	Czechoslovakia	LV	Latvia	TG	Togo
CZ	Czech Republic	MC	Monaco	TJ	Tajikistan
DE	Germany	MD	Republic of Moldova	TT	Trinidad and Tobago
DK	Denmark	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	US	United States of America
FI	Finland	MN	Mongolia	UZ	Uzbekistan
FR	France			VN	Viet Nam
GA	Gabon				

PIPELINED MICROINSTRUCTION APPARATUS
AND METHODS WITH BRANCH PREDICTION
AND SPECULATIVE STATE CHANGING

BACKGROUND AND SUMMARY OF THE INVENTION

The present invention relates generally to improved apparatus and methods for controlling the operation of digital data processors.

5 In the above mentioned concurrently filed application Serial No. 08/174,856, a highly advantageous high speed microinstruction execution architecture is disclosed which is preferably implemented on a single integrated circuit chip or module and which is capable of
10 employing CISC and RISC instructions with multistage pipelining and branch prediction. The present invention provides a very significant and novel enhancement of the approach disclosed in this application which is likewise preferably implemented on a single integrated circuit chip
15 or module and which is also applicable to a wide variety of other architectures.

 The novel enhancement provided by the present invention adds the very important capability of permitting speculative changes of state to occur during execution of
20 a predicted instruction before it is determined whether or not the instruction was correctly predicted. As is well known by those skilled in the art, permitting changes of state for a possibly incorrect instruction is considered to be undesirable, because of the problems associated with

having to undo the incorrect state changes, which is an even more difficult problem where a multistage architecture is employed, as in the aforementioned concurrently filed patent application.

5 In accordance with the present invention, methods and apparatus are provided which permit speculative changes of state to be made during execution of a predicted instruction before the correctness of the prediction has been determined, while also providing for correcting the
10 incorrectly changed state in a highly advantageous and expeditious manner if it is determined that the invention was mispredicted.

 The specific nature of the invention as well as other objects, features, advantages and uses thereof will
15 become evident from the following detailed description of a preferred embodiment in conjunction with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block and schematic diagram of a prior art digital data processing system.

Fig. 2 illustrates the format of a control word used in the system of Fig. 1.

Fig. 3 is a block and schematic diagram illustrating the preferred embodiment disclosed in the aforementioned concurrently filed application.

Fig. 4 graphically illustrates the accessing of microinstructions in response to microinstruction addresses.

Fig. 5 graphically illustrates an example of intermixed CISC and RISC instructions which may be provided in the system of Fig. 3.

Fig. 6 illustrates a typical microinstruction stream which may be applied to the system of Fig. 3.

Fig. 7 is a table illustrating microinstruction flow in the system of Fig. 3 for the microinstruction stream of Fig. 6.

Fig. 8 is a block and schematic diagram broadly illustrating a preferred embodiment of a digital data processing system employing speculative state changing in accordance with the invention.

Fig. 9 is a block and schematic diagram more specifically illustrating the system of Fig. 8.

Fig. 10 illustrates a typical microinstruction stream which may be applied to the system of Fig. 9.

5 Fig. 11 is a table illustrating microinstruction flow in the system of Fig. 10 for the microinstruction stream of Fig. 10.

DETAILED DESCRIPTION

Like numerals and characters refer to like
10 elements throughout the figures of the drawings.

A preferred embodiment of the present invention will be described as applied to the multistage high speed microinstruction execution architecture disclosed in the aforementioned concurrently filed application Serial
15 No. _____. Accordingly, the disclosure in this application will initially be presented with reference to Figs. 1-7.

Reference is initially directed to Fig. 1 which generally illustrates a prior art digital data processing
20 system. For the sake of simplicity, block 8 in Fig. 1 is provided to represent conventional portions of a digital data processing system including a memory storage unit (MSU) 10, an arithmetic logic unit (ALU) 11, an

input/output (I/O) section 12, and a timing control 13, each of which are provided in any of a variety of forms well known in the art.

Typically, synchronous timing is provided for the system of Fig. 1 by the timing control 13 with one microinstruction being performed during each machine cycle. A machine cycle typically corresponds to a single clock time, with storage devices representative of machine state being updated by a clock occurring at the end of each cycle. In order not to confuse the drawings, the particular timing signals fed to the various components of the system in Fig. 1 are not shown, but should be assumed to be provided in a conventional manner.

Continuing with the overall description of Fig. 1, it will be understood that a programmable microinstruction memory 14 is provided for storing a plurality of microinstruction control words CW. A typical control word CW is illustrated in Fig. 2 and will be seen to be comprised of a microinstruction M along with a next address field NA and a condition select field CS. A selected microinstruction control word CW is fetched from microinstruction memory 14 into a microinstruction register 16 in response to a microinstruction address M_{addr} in microinstruction address register 18. The address M_{addr} is

provided during the previous cycle by next address logic 20 in response to selected condition signals 22a from condition select logic 22, next address field NA in microinstruction register 16, and alternate addresses A_{addr} provided by block 8. These alternate addresses A_{addr} permit selection of other next microinstruction addresses, such as may be required for subroutine return or interrupt purposes. The selected condition signals 22a are selected by condition select logic 22 in response to condition select field CS in microinstruction register 16 and microinstruction results M_e from block 8.

During system operation in Fig. 1, the following overlapped (pipelined) executing and fetching operations occur during each machine cycle: (1) the system executes the microinstruction M portion of the control word CW stored in a microinstruction register 16 (this control word having been fetched from the microinstruction memory 14 and stored in the microinstruction memory 16 during the previous cycle), and (2) in response to microinstruction address M_{addr} in address register 18, the system fetches a next microinstruction control word CW from the microinstruction memory 14 and stores it in the microinstruction register 16 for use during the next cycle.

It will be understood with respect to the prior art system of Fig. 1 that, when a microinstruction is unconditional (i.e., there is only one next microinstruction), the fetching operation ((2) above) will
5 always fetch the correct next microinstruction into microinstruction register 16. For such unconditional microinstructions, the pipelined microinstruction fetching and executing operations will be effective to increase the speed of microinstruction flow. However, if the
10 microinstruction is conditional (i.e., there is more than one possible next microinstruction), a predicted next microinstruction address M_{addr} is provided to address register 18 by block 8 in Fig. 1, which could result in fetching the wrong next microinstruction. In such case,
15 steps have to be taken to correct the microinstruction flow.

Various ways are known for correcting for an incorrectly fetched microinstruction as described, for example, in the aforementioned U.S. Patent Nos. 4,155,120
20 and 4,430,706. The disclosures of these patents are to be considered as incorporated herein. The approach taken by these disclosures is to test each predicted microinstruction to determine whether it was correctly predicted. If not, the incorrectly fetched

microinstruction is aborted (i.e., prevented from causing any change in machine state), and provision is made for fetching the correct next microinstruction.

A preferred embodiment of the multistage
5 pipelined, microprogramming architecture disclosed in the
aforementioned concurrently filed application is generally
illustrated in Fig. 3. As shown in Fig. 3, the fetching
and execution of microinstructions are performed in a
pipelined manner using multiple stages. Registers 10, 12,
10 14 and 16 provide for storage between stages. Although not
required, it is of particular advantage to provide clocking
for these multiple stages at a rate which is a multiple of
the basic external clock rate, whereby throughput speed is
increased accordingly. This multiple clock rate may, for
15 example, be at the high rate typically employed in RISC
systems. Such multistage high-speed pipelining as shown in
Fig. 3 is referred to herein as super-pipelining.

Considering Fig. 3 more specifically, four stages
are illustrated, designated as stage 0, stage 1, stage 2
20 and stage 3. Each stage requires only a single clock. In
order not to overly complicate Fig. 3, clock signals are
not illustrated, but should be assumed to be provided in a
conventional manner.

It will be understood by reference to Fig. 3 that stage 0 involves reading out a microinstruction M from a microinstruction memory 100 into a microinstruction register 104 in response to a microinstruction address in microinstruction address register 102. Microinstruction memory 100 may, for example, be an on-chip instruction cache, such as disclosed in the commonly assigned copending patent applications Serial Nos. 08/159,883 and 08/159,683 filed 11/30/93.

During stage 0 in Fig. 3, next address data NA is also read out from microinstruction memory 100 along with microinstruction M. This next address data NA is applied to next address logic 106 which provides a next microinstruction address A_m to microinstruction address register 102 for selecting the next microinstruction to be used during the next clock cycle. Thus, when stage 0 completes, a microinstruction M will have been set up in microinstruction register 104 and the address of the next microinstruction will have been set up in microinstruction address register 102.

In the preferred embodiment being described, a static next microinstruction prediction approach is employed, but it is to be understood that dynamic prediction could also be employed. Accordingly, for each

conditional microinstruction, the next microinstruction is known and is fetched in the same way as an unconditional instruction. For example, if static prediction is employed, the next instruction following a conditional instruction is predetermined by the compiler and embedded in the microcode. Fetching of a next microinstruction thus does not have to wait for testing of the microinstruction execution results from the current or any previously executed microinstruction. Instead, the predicted next microinstruction address A_m is set up in microinstruction address register 102 during the same clock that a current microinstruction M is being read into microinstruction register 104. This operation is exemplified in Fig. 4, which is a graphical representation illustrating the manner in which successive next microinstruction addresses A_{m1} , A_{m2} , A_{m3} in microinstruction address register 102 during respective clock periods 1, 2, 3 cause reading out of corresponding microinstructions $M1$, $M2$, $M3$ in the respective next following clock periods 2, 3, 4.

The predicted next microinstruction A_m set up in microinstruction address register 102 during each clock is provided by next address logic 106 in response to NA read out from microinstruction memory 100. NA may itself be a microinstruction address, or may indicate that the next

sequential (+1) microinstruction address should be used, or may indicate which of a plurality of input addresses A_I should be used as the next microinstruction address A_m . These input addresses A_I are provided for handling
5 subroutine returns, jumps, interrupts, and also provide for ending a microinstruction (when its last microinstruction has been completed) and proceeding to the first microinstruction of the next macroinstruction. Next macroinstructions may also be statically predicted in the
10 preferred embodiment, which is readily accomplished since the flow from one macroinstruction to the next is transparent at the microinstruction level regardless of the number of microinstructions in the macroinstruction.

Fig. 5 graphically illustrates an example of the
15 intermixing of CISC and RISC microinstructions during clock periods 1-10. As shown in Fig. 5, microinstructions M1, M2, M3, M4 constitute a first CISC macroinstruction MCISC-1, the next following microinstructions M5, M6, M7, constitute three RISC macroinstructions MRISC-1, MRISC-2,
20 MRISC-3, and the next three microinstructions M8, M9, M10 constitute a second CISC macroinstruction MCISC-2. The string of RISC macroinstructions MRISC-1, MRISC-2, MRISC-3 intermixed with CISC macroinstructions MCISC-1 and MCISC-2 permits advantage to be taken of RISC performance without

giving up CISC advantages. For example, such strings of RISC instructions can be used to provide RISC operations for those portions of the microcode flow which will result in higher performance.

5 Having described how stage 0 provides for fetching a microinstruction into microinstruction register 104 and for setting up a next microinstruction address A_m in microinstruction address register 102, operations during stage 1 will next be described with reference to Fig. 3.

10 During stage 1, stage 1 control signals 104a derived from microinstruction register 104 cause selected operands to be accessed from a register file 108 and stored in operand registers 110.

 The register file 108 stores the "state" of the

15 virtual machine that is implemented through the hardware/microcode combination. For example, register file 108 stores various addresses (such as A_1 applied to next address logic 106), intermediate results, and other data required for system operation. The register file 108 may

20 comprise an on-chip cache, such as disclosed in the aforementioned commonly assigned copending patent applications S.N. 08/159,883 and 08/159,683. Register file 108 may also interface with an off-chip main memory

(not shown) for receiving data therefrom and for sending data thereto.

As illustrated in Fig. 3, the following control functions are also performed during stage 1: a) Stage 1 control signals 104b from microinstruction register 104 are transferred to a stage 2 control register 112 for use in controlling operations during stage 2; and b) Stage 1 condition select signals CS from microinstruction register 104 are applied to stage 1 condition select logic 114 for selecting particular stage 1 condition signals to be transferred to a stage 2 conditions register 116, also for use in stage 2. Having described how a microinstruction is accessed in stage 0, and how operands are selected in response thereto in stage 1, this description will now proceed to a consideration of stage 2 during which rotation, masking and arithmetic and logic operations are performed on these selected operands. More specifically, during stage 3, the selected operands in operand registers 110 are operated on by rotate logic 118, mask logic 120, and arithmetic and control logic 122 (ALU), which in turn are controlled by stage 2 control signals from stage 2 control register 112. The results of these operations during stage 3 are stored in a microinstruction results register 124. These results in register 124 include flags

125 which are used during stage 3 for determining whether the microinstruction fetched in stage 0 was correctly predicted, as will next be described. Similar to stage 3 , stage 2 control signals from stage 2 control register 112 are transferred to a stage 3 control register 126 for use in controlling operations during stage 3; likewise, stage 2 control signals CS_2 from stage 2 control register 112 are applied to stage 2 condition select logic 128 to select particular conditions to be transferred to a stage 3 conditions register 130, also for use during stage 3.

Execution of a microinstruction is completed during stage 3. During stage 3, the results stored in microinstruction results register 24 in Fig. 3 are written into register file 108. Also, during stage 3, prediction test logic 132 determines whether the microinstruction was properly predicted. This is accomplished by applying flag outputs 125a to prediction test logic 132 along with conditions 130a from stage 3 condition register 130, predicted conditions 131 derived from register file 108, and/or outputs from various logic operations being concurrently performed, including these occurring in other stages. Based on these inputs, prediction test logic 132 determines whether a microinstruction was wrongly predicted and, if so, produces abort signals 132a. These abort

signals 132a initiate the following operations: a) abort signals 132a are applied to register file 108 and other storage and logic to prevent the mispredicted microinstruction from improperly changing system state and/or other system storage: b) abort signals 132a cause the pipeline to be emptied by invalidating the effect of microinstructions fetched following the mispredicted microinstruction. Accordingly, abort signals are used to invalidate the data in address, operand and result registers 102, 110 and 124, in microinstruction and control registers 104, 112 and 126, and in condition registers 116 and 130. This invalidating in response to abort signals 132a may, for example, be accomplished as disclosed in the aforementioned patent nos. 4,155,120 and 4,430,706 and in the aforementioned patent applications S.N. 08/159,883 and S.N. 08/159,683. A preferred manner is to include a validity bit in each involved register, which is settable by the abort signals 132a to indicate that the data in the register is invalid. In order not to confuse Fig. 3, the connections of the abort signals 130a to the various registers are not shown, but may be readily provided by those skilled in the art.

Stage 3 also provides for fetching the correct microinstruction when prediction test logic 130 determines

that a microinstruction was mispredicted. This is accomplished by applying signals 130b from prediction test logic 130 to next address logic 106 in Fig. 3, which causes next address logic 106 to set up the correct
5 microinstruction address (corresponding to the correct microinstruction which should have been fetched) in microinstruction address register 102. Operations then proceed as previously described for stages 0, 1, 2 and 3 in connection with Fig. 3. Since discovery of the
10 mispredicted instruction caused the pipeline to be emptied, operations in invalidated stages will be ignored.

It will be understood that emptying of the pipeline as a result of a mispredicted conditional microinstruction, as described above, can severely affect
15 performance, if it occurs too often, particularly where a multistage architecture is employed having a deep pipeline. However, it has been discovered that conditional branches can be predicted with high accuracy at the microinstruction level, even with static prediction, so that the overall
20 performance of the multistage super-pipelined architecture illustrated in Fig. 3 is able to provide a very significant performance improvement.

The above described multistage super-pipelined architecture offers further advantages. For example, note

in Fig. 3, that additional output signals 122b and 122c are provided from ALU 122, and additional output signals 124b and 124c are provided from results register 124. These output signals may advantageously be applied to logic and storage in other stages to enhance performance. For example, note that outputs 122b and 124b from ALU 122 and results register 124, respectively, are applied as inputs to the operand registers 112. Also, outputs 122c and 124c from ALU 122 and results register 124, respectively, may be provided for use as inputs to microinstruction register 104, control registers 112 and 126, and condition select logic 114 and 128 for affecting the operation thereof in their respective stages. Such interconnections between stages made possible by the pipelined multistage architecture provides greater versatility and flexibility in enhancing operations in each stage, and thus enhances overall microinstruction execution. Another significant advantage is that logical operations which cannot be performed in one clock can be partially performed in a first stage and then completed in the next stage. Furthermore, since multiple stages are available for performing logic related to determining whether a conditional microinstruction was correctly predicted, the correctness of the prediction can be determined in stage 3

without requiring an additional delay for logic testing operations.

Reference is next directed to Figs. 6 and 7 which will be used to illustrate typical operation of the multistage super-pipelined architecture in Fig. 3.

Fig. 6 is an example of a typical microinstruction stream, wherein microinstructions B1 and F1 are conditional, and wherein it is predicted that B2 follows B1 and F2 follows F1.

The table in Fig. 7 illustrates microinstruction flow through stages 0, 1, 2 and 3 for ten clocks periods based on the microcode flow of Fig. 6, wherein the prediction of flow from B1 to B2 is found to be incorrect, while the prediction of flow from F1 to F2 is found to be corrected. It is to be understood that a clock period in Fig. 8 preferably corresponds to a single clock time with storage devices representative of machine state being updated at the end of each clock period.

It will be remembered with reference to Figs. 3 and 4 that: during stage 0, a microinstruction is fetched and stored in microinstruction register 104; during stage 1, operands for the fetched microinstruction are selected and stored in microinstruction operand registers 110; during stage 2, rotate, mask and ALU operations are

performed on the selected operands and the results (and flags) stored in results register 124; and, during stage 3, the results are written into register file 108 and the correctness of a predicted microinstruction is tested. If
5 it is determined during stage 3 that an incorrect microinstruction was fetched, the microinstruction is aborted, the pipeline is emptied, and the address of the correct microinstruction is set up in address register 102.

Initially, clock periods 1-4 in Fig. 7, which do
10 not involve abortion, will be considered in more detail. As shown in Fig. 7, during clock periods 1-4, stage 0 fetches respective microinstructions A1, B1, C1 and D1; during clock periods 2-4, stage 1 selects operands for respective microinstructions A1, B1 and C1; during clock
15 periods 3-4, stage 2 performs mask, rotate and ALU operations for respective microinstructions A1 and B1; and during clock period 4, stage 3 performs result storage and next microinstruction address selection for microinstruction A1. Since no incorrect prediction is
20 discovered during clock periods 1-4, no abortion occurs.

During clock period 5, stages 0, 1, 2 and 3 perform their respective operations on microinstructions E1, D1, C1 and B1, as indicated in Fig. 7. However, during stage 3 of clock 5, the prediction test logic 130 will

determine that the microinstruction C1 was incorrectly predicted, and that microinstruction B2 should have been fetched instead. As a result, an abortion occurs (as previously described), which prevents writing in register
5 file 108, empties the pipeline by invalidating the data in registers 104, 110, 112, 116, 124, 126 and 130, and causes the correct microinstruction address (i.e., the address of B2) to be set up in microinstruction address register 102.

Thus, as illustrated for clock period 6 in
10 Fig. 7, during stage 0 of clock period 6, the correct microinstruction B2 is fetched into microinstruction register 104. Note that an "X" is marked in stages 1, 2 and 3 of clock period 6 to indicate that operations of these invalidated (emptied) stages are ignored during clock
15 period 6.

During clock period 7 in Fig. 7, microinstruction F1 is fetched during stage 0, and operands are selected for microinstruction B2 during stage 1, as shown. Stages 2 and 3 continue to be ignored during clock 7 (as indicated
20 by the "X" provided for these stages in clock 7), since valid data corresponding to the correct microinstruction B2 has not yet propagated thereto.

In clock period 8 in Fig. 7, microinstruction G1 is fetched in stage 0. Note in Fig. 6 that F1 is a

conditional microinstruction, in which case microinstruction G1 is fetched based on the assumed prediction. Also, during clock 8, operands are selected for microinstruction F1 in stage 1, and mask, rotation and
5 ALU operations are performed for microinstruction B2 in stage 2. Stage 3 of clock 8 is again marked with an "X" since valid data still has not propagated thereto.

In clock period 9, the pipeline is again filled, since all four stages 0, 1, 2 and 3 are now valid. In
10 clock period 9, the next microinstruction G2 is fetched during stage 0, operands are selected for microinstruction F1 during stage 2, and execution is completed for microinstruction B2 in stage 3.

In clock period 10, which is the last clock
15 period illustrated in Fig. 7, the next microinstruction G3 is fetched during stage 0, operands are selected for microinstruction G2 in stage 1, rotate, mask and ALU operations are performed for microinstruction G1 in stage 2, and microinstruction completion operations are
20 performed for microinstruction F1 in stage 3. Since microinstruction F1 is conditional, as shown in Fig. 6, stage 3 also determines whether the selection of the predicted microinstruction G1 was correct. Since the prediction will be found to be correct, as assumed for this

example, no abortion occurs and the flow continues without interruption.

Having described with reference to Figs. 1-7 the microinstruction execution architecture disclosed in the
5 aforementioned concurrently filed patent application Serial
No. _____, a preferred manner for incorporating the
present invention therein will next be described with
reference to Figs. 8-11. First, however, it will be helpful
in understanding the present invention to consider some of
10 the reasons as to why performance can be degraded when
state changes are delayed until it is determined that an
instruction has been correctly predicted.

It will be understood that several data hazards
exist when employing pipelined instruction execution
15 architectures that allow concurrent or out-of-order
execution, as is typical in present day systems. As is
well known, a data hazard refers to a situation where it is
possible to read or write incorrect data, usually due to
overlapping instruction executions. One example of a data
20 hazard is a read-after-write (RAW) hazard, where an
instruction "writes" data that a subsequent instruction
must "read", in which case the "write" must be performed
before the "read" is allowed to occur. Another example of
a data hazard is a write-after-read (WAR) hazard, where an

instruction "reads" data that a subsequent instruction must "write", in which case the "write" must occur after the "read". An additional example of a data hazard is a writeafter-write (WAW) hazard, where an instruction writes data that a subsequent instruction must also write, in which case the "writes" must occur in the proper order.

Pipelined architectures typically solve data hazards, such as exemplified above, by forcing instructions to write their respective results (changing machine state) in a predetermined order. Such forcing of instruction writing can cause significant delays in the speed of instruction execution, particularly for pipelined machines employing branch prediction, since further delays are introduced by having to delay the writing of results until it is determined that an instruction was correctly predicted.

Fig. 8 is a block and schematic diagram broadly illustrating a preferred embodiment employing speculative state changing in accordance with the invention.

It is to be understood with respect to Fig. 8 that a microinstruction architecture similar to that described in Fig. 3 is preferably employed. More specifically, Fig. 8 shows three stages of a pipelined multistage architecture employing intermediate registers

I1, I2 and I3 which receive instruction execution results from execution stages EX-1, EX-2 and EX-3, respectively. Intermediate registers I1, I2 and I3 may, for example, correspond to registers 104, 110 and 124 in Fig. 3 respectively, while execution stages EX-1, EX-2 and EX-3 may correspond to stages 1, 2 and 3 in Fig. 3. Machine state block 200 shown in Fig. 8 corresponds to both the register file 108 and the other machine state 109 in Fig. 3. Also, for greater clarity, various portions of Fig. 3 are omitted from Fig. 8, but may be provided in the same manner as shown in Fig. 3.

Execution results applied to machine state 200 in Fig. 8 from execution stages EX-1, EX-2 and EX-3 are indicated as E1, E2, and E3, respectively, while the particular machine state values applied to stages EX-1, EX-2 and EX-3 for use thereby are indicated as S1, S2 and S3, respectively.

As shown in Fig. 8, restore registers 202 communicate with machine state 200 via lines 201. These restore registers 202 are provided to store backup copies of machine state in order to permit machine state 200 to be changed in response to instruction execution in a stage, even though it has not yet been determined whether the instruction which produced the change correctly predicted.

This is accomplished by allowing early state changes to be made in one or more state values, while storing backup copies (via lines 201) of the original state values in corresponding restore registers 202. Accordingly, if at a later time it is determined that the instruction which caused the state changes was mispredicted, the backup copies in the corresponding restore registers are applied (via lines 201) to machine state 200 to restore the incorrectly changed state values to their original (correct) values.

Since the pipelined architecture of Fig. 8 permits state values to be changed without having to wait until it is determined that the corresponding instruction has been correctly predicted, a significantly higher speed of instruction execution can be achieved. For example, for instructions that can complete quickly and for which instruction execution can be speeded up by changing state as early as possible, allowing early changes of state as described above can significantly improve performance. Since it is desirable to limit the number of restore registers 202 provided, backup should be provided primarily for those state changes which will produce the greatest improvement in performance.

Reference is next directed to Figs. 9-11 which will be used for illustrating a specific example of the operation of the speculative state changing architecture shown in Fig. 8.

5 Fig. 9 is basically similar to Fig. 8, with counter 303 in Fig. 8 being representative of a state value in block 200 for which speculative state changing can advantageously be provided. Similarly, EX-2 and EX-3 restore registers 305 and 307 in Fig. 9 are representative
10 of restore registers 202 in Fig. 8 which in Fig. 9 are used for providing backup of counter 303. Execution blocks EX-1, EX-2 and EX-3 and intermediate registers I1, I2 and I3 in Fig. 10 are as previously described in connection with Fig. 8.

15 For the purposes of this operative example being described, it will be assumed that the microinstruction flow of Fig. 10 is presented, wherein A is a conditional microinstruction and has two possible next microinstructions, microinstruction B or D. Also, it will
20 be assumed that each of microinstructions A, B, D and F in Fig. 10 increment counter 303 in Fig. 9 during stage EX-1, and that counter 303 is not otherwise incremented. This is indicated in Fig. 9 by stage EX-1 applying a new counter value CV to counter 303 via multiplexor 310 which is what

normally occurs when no misprediction is detected during EX-3, as will hereinafter be explained. When CV increments counter 303 during stage EX-1, a backup copy of the original (unchanged) count value is stored in EX-2 restore register 305. Note that the counter output 303a is also applied to execution block EX-1 for use thereby during each clock period. Also note, as shown in Fig. 9, that EX-2 restore register 305 transfers its backup value to EX-3 restore register 307 at the end of each microinstruction, so as to permit its use in stage EX-3 when its respective microinstruction is executed thereby.

Execution block EX-3 in Fig. 9, similar to stage 3 in Fig. 2, determines whether a microinstruction was mispredicted. If so, EX-3 provides appropriate outputs Ab for performing abort operations, as previously described for stage 3 in Fig. 3. In addition, when a misprediction is detected, EX-3 in Fig. 9 provides a restore counter output Rc which is applied to multiplexer 310 for selecting the count value in EX-3 restore register 307 to be applied to counter 303 instead of the counter value CV from EX-1. This, in effect, restores counter 303 to the value that it had prior to being changed during EX-1 by the microinstruction which was found to have been mispredicted. If EX-3 determines that the microinstruction was correctly

predicted, counter 303 is not restored but is set in the usual manner in accordance with the new counter value CV provided by EX-1.

Next to be described with additional reference to Fig. 11, is the specific manner in which the exemplary microinstruction flow shown in Fig. 10 propagates through the speculative execution architecture of Fig. 10 during clock periods 1-7. It will be assumed for this example that microinstruction B is predicted to follow conditional microinstruction A, but that this is a misprediction, since microinstruction A should have been followed by microinstruction D. It will additionally be assumed that counter 303 initially has a "1" count value.

Note that Fig. 11 indicates which microinstruction is being operated on by each of stages EX-1, EX-2 and EX-3 during each of the illustrated clock periods 1-7. As in Fig. 8, a clock period in Fig. 12 preferably corresponds to a single clock time with storage devices representative of machine state being updated at the end of each clock period. Fig. 11 also indicates the values of counter 303 and counter restore registers 305 and 307 during each clock period. Operation during each clock will now be considered in more detail.

During clock period 1 in Fig. 11, microinstruction A is operated on by stage EX-1, as shown. Also, counter 303 will be seen to have the assumed initial count value of "1". Activities which are not pertinent
5 during clock period 1 are indicated by a "?".

During clock period 2 in Fig. 11, microinstruction A has propagated to stage EX-2, and the predicted next microinstruction B is in stage EX-1. Also, during clock period 2, counter 303 has a "2" count value as
10 a result of having received an incremented CV value of "2" from stage EX-1 operated on microinstruction A during clock period 1. Also, EX-1 counter restore register 305 has a "1" value during clock period 2 as a result of the original "1" value of counter 303 having been stored therein during
15 clock period 1 to provide a backup copy.

As shown for clock period 3 in Fig. 11, microinstruction C1 is in stage EX-1, microinstruction B is in stage EX-2, and microinstruction A is in stage EX-3. Since it is assumed that the CV value provided by EX-1
20 during microinstruction B (like microinstruction A) causes counter 303 to be incremented during EX-1, counter 303 will have a "3" count value during clock period 3, while EX-2 restore register 305 will have a "2" value corresponding to the prior "2" count value of counter 303.

Also, the EX-2 counter 307 has a "1" value during clock period 3 as a result of the previous "1" count in counter restore register 303 having been transferred thereto during clock period 2, as indicated in Fig. 10.

5 Since microinstruction A is in stage EX-3 during clock period 3, the correctness of the predicted branch from microinstruction A to microinstruction B will be tested. As mentioned previously, for the example being presented, it is assumed that stage EX-3 will determine
10 that microinstruction B was incorrectly predicted, and that microinstruction D should have followed microinstruction A.

 As a result of determining during clock period 3 that microinstruction B was mispredicted, during clock
15 period 4 the correct microinstruction D is caused to be applied to stage EX-1 (as described for example in connection with stage 3 of Fig. 3), while the other stages EX-2 and EX-3 as well as the EX-2 and EX-3 restore registers 305 and 307 are invalidated, as indicated by "X".
20 Note that, during clock period 4, counter 303 will have been restored to its original "1" value (i.e., the value it had prior to execution of microinstruction B in clock period 1) which restoration occurs as a result of restore output Rc from EX-3 during clock period 3 causing

multiplexor 310 to select the original "1" count value in EX-3 counter restore register 307 for application to counter 303.

As shown in Fig. 11, during the next following
5 clocks 5-7 in Fig. 11, operations on the corrected microinstruction flow D, E, F, etc. proceed as shown in Fig. 11, with counter 303 and EX-2 and EX-3 restore registers 305 and 307 providing backup copies for counter 303, as previously described. The resulting completed
10 microinstruction flow for clocks 1-7 is A, D, E, as illustrated by the stage EX-3 line in Fig. 11.

It is to be understood that the present invention is not limited to the preferred embodiment of the invention described herein, since many variations in construction
15 arrangement, use and operation are possible within the scope of the invention.

Accordingly, the present invention is to be considered as encompassing all modifications, variations and adaptations coming within the scope of the invention as
20 defined by the appended claims.

What is claimed is:

1. In a data processing system, the combination comprising:

an addressable memory for storing instructions,
at least particular ones of said instruction in said memory
5 being conditional instructions having a plurality of
possible next instructions;

addressing means for selectively addressing said
memory for reading out a selected instruction therefrom;

storage means for storing execution results and
10 other machine state; and

executing means for executing instructions read
out from said memory and for changing said storage means in
response thereto;

said executing means including multiple stages
15 operating in a pipelined manner on instructions read out
from said memory, wherein execution by at least one of said
stages provides a next instruction address for said
addressing means which is a predicted next instruction for
a conditional microinstruction;

20 said executing means providing for detecting and
correcting an incorrect instruction prediction while also
permitting said storage means to be changed in response to
an instruction before detection of the correctness thereof.

What is claimed is:

2. Pipelined microinstruction execution apparatus comprising:

an addressable memory for storing microinstructions, said microinstructions instituting CISC and RISC instructions, a CISC instruction comprising a plurality of microinstructions stored in said memory and a RISC instruction comprising a single microinstruction stored in said memory, at least particular ones of said microinstructions in said memory being conditional microinstructions having a plurality of possible next microinstructions;

addressing means for selectively addressing said memory for reading out a selected microinstruction therefrom; and

15 executing means for executing a microinstruction read out from said memory by said addressing means;

said executing means including multiple stages operating in a pipelined manner on microinstructions read out from said memory, said executing means providing for communication of control, condition and data signals among said stages for use thereby during execution wherein execution by at least one of said stages provides a next

microinstruction address for said addressing means which is
a predicted next microinstruction for a conditional
25 microinstruction,

said executing means also providing for detecting
and correcting an incorrect microinstruction prediction.

1/7





FIG. 2
(PRIOR ART)

FIG. 3




FIG. 4

CLOCK PERIOD	1	2	3	4
MICROINSTRUCTION ADDRESS A_m	A_{m1}	A_{m2}	A_{m3}	--
MICROINSTRUCTION M	--	M1	M2	M3

CLOCK PERIOD	1	2	3	4	5	6	7	8	9	10
MICROINSTRUCTIONS	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10

MCISC-1

MRISC-1


MRISC-2

MCISC-2

MRISC-3

FIG. 5


4/7

**FIG. 6**

CLOCK PERIOD	1	2	3	4	5	6	7	8	9	10
STAGE 0	A1	B1	C1	D1	E1	B2	F1	G1	G2	G3
STAGE 1		A1	B1	C1	D1	X	B2	F1	G1	G2
STAGE 2			A1	B1	C1	X	X	B2	F1	G1
STAGE 3				A1	B1	X	X	X	B2	F1

FIG. 7

5/7

**FIG. 8**

6/7




FIG. 9

7/7




FIG. 10

CLOCK PERIOD	1	2	3	4	5	6	7
EX-1	A	B	C1	D	E	F	G
EX-2	?	A	B	X	D	E	F
EX-3	?	?	A	X	X	D	E
COUNTER 303	1	2	3	1	1	2	3
RESTORE REG 305	?	1	2	X	X	1	2
RESTORE REG 307	?	?	1	X	X	1	1

FIG. 11

INTERNATIONAL SEARCH REPORT

Intern. al Application No
PCT/US 94/14929

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 G06F9/28

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.
Y	US,A,4 430 706 (SAND DUANE B) 7 February 1984	1
X	see the whole document ---	2
Y	IBM TECHNICAL DISCLOSURE BULLETIN, vol. 29, no. 3, August 1986 NEW YORK, US, pages 991-993, 'Use of a second set of general purpose registers to allow changing general-purpose registers during conditional branch resolution' see the whole document ---	1
A	US,A,4 155 120 (KEEFER DAVID E ET AL) 15 May 1979 see the whole document -----	1,2

☐ 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

18 April 1995

Date of mailing of the international search report

10.05.95

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

Klocke, L

INTERNATIONAL SEARCH REPORT

information on patent family members

Intern. Application No

PCT/US 94/14929

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US-A-4430706	07-02-84	NONE	
US-A-4155120	15-05-79	NONE	