


METHOD AND APPARATUS FOR AVOIDING WRITEBACK CONFLICTS BETWEEN
EXECUTION UNITS SHARING A COMMON WRITEBACK PATH



BACKGROUND OF THE INVENTION

Field of the Invention

The present invention relates to the field of
5 microprocessor architecture. More specifically, the present
invention relates to the field of microprocessor architecture
for coordinating instructions executed in parallel which
require common resources by providing a method and apparatus
for extending execution unit pipelines.

10

Art Background

As the computer revolution has progressed the quest
of microprocessor developers has been to develop chips
exhibiting more power and faster performance. Initial
15 efforts focused essentially on increasing transistor
populations on single microprocessor integrated circuits.
That effort continues with today's microprocessors now
housing literally millions of transistors on a single chip.
Further integration has allowed processor clock speeds to be
20 greatly increased with the increased density of transistors.

In addition to squeezing performance by overcoming
physical limitations, microprocessor design has developed
into an art form. Microprocessors are divided into discrete

functional blocks through which instructions are propagated one stage at a time. This allows for pipelining of instructions such that when one instruction has completed the first stage of processing and moves on to the second stage, a
5 second instruction may begin the first stage. Thus, even where each instruction requires a number of clock cycles to complete all stages of processing, pipelining provides for the completion of instructions on every clock cycle. This single-cycle throughput of a pipelined microprocessor greatly
10 increases the overall performance of computer systems.

Other enhancements to microprocessor design include the development of superscalar microprocessors which are capable of launching more than one instruction at the initial stage of the pipeline per clock cycle. Likewise, in a
15 superscalar microprocessor, frequently more than one instruction completes on each given clock cycle. Other development efforts have gone into the simplification of microprocessor instruction sets, developing reduced instruction set (RISC) microprocessors which exploit the fact
20 that many simple instructions are more commonly executed than some complicated instructions. Eliminating the complicated instructions from the instruction set provides for a faster executing pipeline. Complicated instructions are carried out by combinations of the more simple instructions.

25 In order for pipelined microprocessors to operate efficiently, an instruction fetch unit at the head of the

pipeline must continually provide the pipeline with a stream of instructions. However, conditional branch instructions within an instruction stream prevent an instruction fetch unit at the head of a pipeline from fetching the correct
5 instruction until the condition is resolved. Since the condition will not be resolved until further down the pipeline, the instruction fetch unit may not be able to fetch proper instructions.

To overcome this problem, many pipelined
10 microprocessors use branch prediction mechanisms that predict the outcome of branches and then fetch subsequent instructions according to branch prediction. Branch prediction is achieved using a branch target buffer to store the history of a branch instruction based upon the
15 instruction pointer or address of that instruction. Every time a branch instruction is fetched, the branch target buffer predicts the target address of the branch using the branch history. The use of branch prediction in a pipelined instruction architecture is often referred to as speculative
20 execution.

In addition to speculative execution, substantial increases in instruction throughput are achievable by implementing out-of-order dispatch of instructions to the execution units. Many experiments have confirmed that
25 typical von Neumann code provides substantial parallelism and hence a potential performance boost by use of out-of-order

execution. Out-of-order execution is possible when a given instruction does not depend on previous instructions for a result before executing. With out-of-order execution, any number of instructions are allowed to be in execution in the
5 execution units, up to the total number of pipeline stages for all the functional units.

In a processor using out-of-order execution, instruction issuing is stalled when there is a conflict for a functional unit or when an issued instruction depends on the
10 result of an instruction that is not yet computed. In order to prevent or mitigate stalls in decoding, previous texts have described the provision of a buffer known as a reservation station (RS) between the decode and execute
15 stages. The processor decodes instructions and places them into the reservation station as long as there is room in the buffer and at the same time examines instructions in the reservation station to find those that can be dispatched to the execution units (that is, instructions for which source
20 operands and execution units are available). Data-ready instructions are issued from the reservation station with little regard for their original program order. For a further description of the use of a reservation station and out-of-order execution, see Mike Johnson, Superscalar Microprocessor Design and Prentice-Hall, Inc., 1991, Chapters
25 3 and 7.

In microprocessors implementing reservation stations, including superscalar processors, the reservation station will generally provide ready instructions directly from the reservation station to an appropriate execution unit. In many cases execution units may be grouped together coupled from a single dispatch port from the reservation station. For example, one implementation may have floating point add and floating point multiply execution units coupled through the same dispatch port. Floating point add and floating point multiply execution units may have differing length internal pipelines. A problem arises because these commonly grouped execution units may share a single writeback path from the execution units. If the reservation station dispatches a first instruction to an execution unit such as to the floating point multiply unit and then subsequently dispatches an instruction to the floating point add execution unit in a common group of execution units, it is possible that multiple execution units will be ready to write back a result at the same time. Because the execution units share a common writeback port, a contention develops.

It is an object of the present invention to overcome the contention of execution units which share a common writeback port and which face this contention possibility.

SUMMARY OF THE INVENTION

From the foregoing it can be appreciated that there is a need to provide a means for overcoming contentions between execution units which share a common writeback path and which may require simultaneous access to the writeback path. Therefore, it is an object of the present invention to provide a means for changing the internal pipeline length of an execution unit which may be involved in a writeback contention with another execution unit on a shared writeback port.

It is also an object of the present invention to be able to detect potential conflicts on a shared writeback port between multiple functional units.

It is another object of the present invention to provide a means of inhibiting or delaying a reservation station from dispatching instructions which may lead to writeback conflicts.

These and other objects of the present invention are provided by a means and method for lengthening the pipeline of functional units that would otherwise be involved in a writeback conflict. Logic circuitry is provided for comparing the depths of two concurrently executing execution unit pipelines to determine if a conflict will develop. When it appears that two execution units will attempt to write

back at the same time, the execution unit having a shorter pipeline will be instructed to add a stage to its pipeline, storing its result in an added register for one clock cycle. After the conflict has been resolved, the instruction to
5 lengthen the pipeline of a given functional unit will be rescinded.

The present invention also provides for a means and method for multi-stage execution units to signal the reservation station to delay the dispatch of various
10 instructions to avoid conflicts between those execution units. For example, a three-stage pipeline execution unit will signal the reservation station not to dispatch any single-cycle instructions that will need to be written back concurrently over a common write-back port. Likewise,
15 execution units requiring an undetermined number of pipeline stages are equipped with means for signaling the reservation station to inhibit dispatch of instructions to execution units sharing a common writeback report a known number of cycles in advance to avoid contentions.

BRIEF DESCRIPTION OF THE DRAWINGS

The objects, features and advantages of the present invention will be apparent from the following detailed description in which:

5 Figure 1 illustrates a basic computer architecture which may incorporate the present invention.

Figure 2 illustrates a block diagram of a portion of a microprocessor incorporating the present invention.

10 Figure 3 illustrates a more detailed figure of the pipeline of two execution units sharing a common writeback port and the means for avoiding conflicting use of the writeback port.

15 Figure 4 illustrates execution units which signal a reservation station in order to inhibit the reservation station from dispatching instructions which will conflict with a need for a common writeback path.

DETAILED DESCRIPTION

The present invention provides a method and apparatus for avoiding writeback conflicts between execution units sharing a common writeback port. In the following
5 description numerous specific details are set forth, such as particular execution unit pipeline lengths and instruction codes, in order to provide a thorough understanding of the present invention. It will be appreciated, however, by one skilled in the art that the present invention may be
10 practiced without such specific details. In other instances, well known control structures and gate level circuits have not been shown in detail in order not to obscure the present invention. Particularly, many functions are described to be carried out by various logic circuits. Those of ordinary
15 skill in the art, having been described the various functions, will be able to implement the necessary logic circuits without undue experimentation.

Referring first to Figure 1, a typical computer system which may utilize a microprocessor that implements the
20 present invention is shown. As shown in Figure 1, there is computer 100 which comprises three major components. The first of these is the input/output (I/O) circuit 101 which is used to communicate information in appropriately structured form to and from other parts of the computer 100 as well as
25 out of the computer 100. Also shown as part of the computer

100 is the central processing unit (CPU) 102 and memory 103. These two latter elements are those typically found in most general purpose computers and almost all special purpose computers. In fact, the several elements contained within
5 computer 100 are intended to be representative of this broad category of data processor. Many commercially available computers having differing capabilities may be utilized with a processor incorporating the present invention.

Also shown in Figure 1 is an input device 105, shown in
10 a typical embodiment as a keyboard. There is also shown as an input device a graphics tablet 107. It should be understood, however, that the input device may actually be in any other well-known input device (including, of course, another computer). A mass memory device 104 is coupled to
15 I/O circuit 101 and provides additional storage capabilities for the computer 100. The mass memory may include other programs, fonts for different characters and the like and may take the form of magnetic or optical disc drive or any other well-known device. It will be appreciated that the data
20 retained within mass memory 104, may, in appropriate cases, be incorporated in standard fashion into computer 100 as part of memory 103.

In addition, three typical computer display devices are illustrated, the display monitor 108, the plotter 109 and a
25 laser printer 110. Each can be used to display images or documents or other data utilized by the computer 100. A

cursor control device 106, such as a mouse, trackball or stylus are also coupled to I/O circuit 101. Other pointing devices may suitably be used as appropriate.

The preferred embodiment of the present invention is intended to be implemented in a microprocessor that executes the well-known Intel Architecture Microprocessor instruction set. The present invention may, however, be implemented in any microprocessor having a microarchitecture in which a plurality of functional units share a common writeback path and which may face the contention of multiple functional units simultaneously needing to write back over the common writeback path.

The microprocessor of the present invention, prior to executing Intel Architecture Microprocessor instructions, decodes them into a more simple, stylized sequence of "micro operations." The micro operations are then analyzed and rescheduled according to resolved dependencies thus implementing out-of-order execution. The sequence below shows the mapping of six instructions of the Dhrystone trace into their constituent micro operations ("Gn" is general register n):

	Intel Architecture Instructions	micro operations
	000014e8: PUSHw BP <TOS fbfffc3c>	STw IN (G5, G4) OUT (Memfbfffc3c) SUB IN (SPO, CO) OUT (SPO)
5	000014e9: MOVw EA, R ESP EBP	MVw IN (G4) OUT (G5)
	000014eb: LEAVE <TOS fbfffc3c>	ADD IN (G5, CO) OUT (G4) LDw IN (G5, Memfbfffc3c) OUT (G5)
	000014ec: RETn	LDw IN (G4, Memfbfffc40) OUT (TO) ADD IN (SPO, CO) OUT (SPO)
10		JMP IN (TO) OUT ()
	00001625: XORw EA, R EDI EDI	XOR IN (G7, G7) OUT (G7)
	00001625: ADDw Ib, EA ESP <IMM 18>	ADD IN (G4) OUT (G4, flags 0)

It can be seen that the PUSH of the BP register is converted into two micro operations. The first performs the store of the register and the second adjusts the stack pointer. These two micro operations are not dependent on each other. The "microcode" corresponding to the PUSH has been written explicitly to allow maximal parallelism wherever possible. It can be seen that some complex instructions translate into multiple micro operations which may or may not be related to one another, and simple instructions map directly into a unique micro operation (such as XOR or ADD).

Figure 2 illustrates a portion of the microprocessor incorporating the present invention. The preferred embodiment microprocessor implements out-of-order execution. The portion of the microprocessor illustrated in Figure 2 can be considered the out-of-order core of the microprocessor. The first element of the out-of-order core illustrated in Figure 2 is the reservation station 20. The reservation station is a buffer in front of the functional units that holds micro operations that are waiting for needed resources in order to execute. These needed resources

include the source data operands (that previous micro operations may not yet have created) and a functional unit with which to execute the micro operation.

The reservation station performs several functions.

5 It serves as a buffer for instructions and for source data that is already available. This source data may have already been stored in the reservation station and be waiting for an instruction to act upon it. In other cases, the reservation station maintains decoded instructions and is "watching" all

10 the result/writeback buses from all the execution interfaces "looking" for source data that it needs for its micro operations. When the watched-for data arrives on the result bus 120, the reservation station grabs it and writes it into the appropriate source data field.

15 Once all the source operands for a micro operation are in the reservation station, the reservation station determines when an appropriate execution unit is available and schedules the data ready micro operation for dispatch to an available execution unit. The reservation station 20 of

20 Figure 2 is illustrated having five dispatch ports. Primarily, micro operations in the embodiment illustrated will be dispatched over either Dispatch Port 0 or Dispatch Port 1.

In the preferred embodiment, there are five

25 execution units coupled to Dispatch Port 0 from reservation station 20. These are the floating point add unit 31, the

floating point and integer multiply unit 32, the floating point divide unit 33, integer execution unit 34, and a wire unit 35. As was described, the complex Intel Architecture Microprocessor instruction set is reduced to a simpler stream
5 of micro operations. These micro operations, when they are data-ready, are operated on by the various execution units. Most applications heavily utilize the integer execution unit. Integer execution unit 34 is designed to be able to process data-ready micro operations in one clock cycle. It is
10 therefore referred to as a single-cycle functional unit. The integer execution unit 34 receives an instruction code sequence and source operand data having a width of 32 bits per operand. The floating point execution units 31, 32 and 33 carry out more complicated tasks when they receive data-
15 ready micro operations. These functional units receive source data operands having a width of 86 bits.

In the preferred embodiment, a floating point add instruction requires three clock cycles to execute, an integer multiply requires four clock cycles, and a floating
20 point multiply instruction requires five clock cycles to execute. Floating point divides and square roots, which are both carried out by the floating point divide unit 33, have variable length latencies depending upon the precision required for their result. The wire unit 35 is a hybrid
25 execution unit implemented in the preferred embodiment for various catch-all instructions. It is similar to a floating point unit in that its input stream takes 86-bit operands,

but it is more like an integer execution unit in that it performs its operations in a single clock cycle.

Dispatch port 1 from reservation station 20 also has two execution units coupled to it. There is an integer execution unit 36 and a jump unit 37. The integer execution unit 36 may be identical to the execution unit 34, or it may have additional or different functionality. Two integer execution units are provided because of the high volume of integer micro operations common in most code thereby enhancing performance of the overall system. Of course, additional integer execution units may be provided in alternative architectures.

After an execution unit executes a micro operation, its result is propagated on a write back bus 120 to the reorder buffer 40 and back to the reservation station 20. The reorder buffer 40 serves as the place where the results of out-of-order operations are collected so that they can later be committed to machine state in strict von Neumann program order. Thus, the reorder buffer is where the out-of-order core reassembles the instruction stream into its original program order. At the same time that execution units write back to the reorder buffer 40, they simultaneously write back to the reservation station 20, because the result of an execution unit's operation may provide a source operand needed for the execution of another micro operation waiting in the reservation station.

As can be seen from Figure 2, the five execution units from Dispatch Port 0 all use the same writeback path 120 to the reorder buffer and to the reservation station 20.

5 This is where certain conflicts may develop and is the motivation for the present invention. A conflict develops where, for example, an integer multiply instruction and floating point add instruction are issued on consecutive clock cycles from the reservation station. Because the

10 integer multiply executes in a four cycle pipeline and the floating point add unit has a three cycle pipeline, consecutively issued instructions to these two units creates a situation in which both desire to provide result data to the write back bus 120 simultaneously. This cannot be done.

15 Similarly, if an integer instruction is dispatched to the integer execution unit 34 two clock cycles after a floating point add has been dispatched or three clock cycles after an integer multiply has been dispatched, the integer execution unit and the multiply unit will also both need to write back

20 on the same clock cycle. This also cannot happen.

Figure 3 illustrates a more detailed view of two execution units coupled to write back to the reorder buffer 40 and reservation station 20 over writeback bus 120. The floating point add unit 31 and multiply unit 32, for example,

25 receive data-ready micro operations from the reservation station. The floating point add unit 31 comprises a three-

stage execution pipeline. Thus, though it requires three clock cycles to complete one micro operation, it may receive a micro operation every clock cycle and provide an output every clock cycle. The three stages of the pipeline are
5 illustrated in Figure 3 as the first stage functional circuit block 41, second stage functional circuit block 42 and third stage functional circuit block 43. At the front end of each stage of the execution unit pipeline are latches or registers for buffering the data stream into each
10 functional stage. These are illustrated as latches 44, 45 and 46, respectively preceding the first, second and third floating point add pipeline stages.

Similarly, the integer portion of multiply functional unit 32 receives its micro operations from the
15 reservation station, but has a four stage execution pipeline shown as first stage 47, second stage 49, third stage 50, and fourth stage 51. In front of each of these stages, respectively, are latches or buffers 52, 53, 55 and 56.

As was described above, when a dispatched integer
20 multiply instruction is followed one clock cycle later by a dispatched floating point add instruction, both functional units will be ready to write back result data at the same time over the same writeback bus 120, a writeback contention. To avoid this, the pipeline of the floating point add
25 functional unit 31 is conditionally extended by one clock cycle. A fourth latch or pipe extension buffer 60 receives

the output from the third pipe stage 43 of the unit.
Buffering the functional unit's result through a fourth
buffer delays, by one clock cycle, the output of the unit.
Thus, in the case where a floating point add has been
5 dispatched one clock cycle after a floating point multiply,
its result will need to be written back one clock cycle after
the floating point multiply result is written by utilizing
the pipe extending register. As an alternative to the pipe
extension buffer 60, other one cycle buffering techniques may
10 be implemented, such as a one cycle delay line.

To control the pipeline length of the floating
point add unit 31, a multiplexing circuit 61 is implemented.
The multiplexer circuitry 61 receives as input signals a
result directly from the third pipe stage 43 of the floating
15 point add unit and a previous clock cycle's buffered result
from the pipeline extending buffer 60. The multiplexer 61 is
controlled by control logic 62 which determines when the
pipeline length of the floating point add unit should be
extended. When the control logic 62 determines that a
20 writeback conflict will exist, it instructs the multiplexer
61 to select the signal coming through the pipe extension
buffer 60 to be the signal to be written back from the
functional unit. When no writeback conflict exists and no
writeback result is pending in the extension buffer 60, the
25 control logic instructs the multiplexer 61 to select the
output directly from the third pipe stage 43 of the floating
point add unit.

In determining whether the floating point add unit should extend the functional unit's pipeline length, the control logic 62 measures the activity occurring at the different pipe stages of the floating point add and integer multiply functional units. The control logic 62 may
5 implement comparing circuitry which determines that if a multiply is in multiply pipe stage 4 at the same time a floating point add is in the third pipe stage of that functional unit, then there will be a writeback conflict, and
10 the floating point add unit needs to extend its pipeline length by one clock cycle. Other comparisons may be performed, so long as the comparison is designed to indicate that a writeback conflict will exist. Those of ordinary skill in the art will be able to implement the control logic
15 62 without undue experimentation for determining when a writeback conflict exists.

When no writeback conflict will exist after one has previously existed, the control logic 62 of the floating point add unit 31 will reset the length of the floating point
20 add pipeline to be three clock cycles. If valid bits are implemented in the buffering elements, then the pipe extension buffer 60 will have its valid bit reset indicating a three clock-cycle pipeline.

The above discussion of lengthening a functional
25 unit pipeline considered the case of a writeback conflict where two functional units have a one pipe stage pipeline

length difference. It is of course possible for two functional units to have a pipeline length difference of more than one pipe stage. In order to optimally implement the present invention for such a case, the pipe extension buffer
5 60 needs to provide buffering for N deep entries where N is the difference in the number of pipeline stages of the conflicting functional units.

Another aspect of the present invention will be described simultaneously with respect to Figures 2, 3 and
10 4. Other writeback conflicts can develop with respect to those execution units receiving their inputs from Dispatch Port 0 of the reservation station 20. There are two single-cycle execution units, the integer execution unit 34 and wire unit 35 (Figure 2). If a micro operation is dispatched to
15 either of these execution units when a floating point unit is at the final stage of its pipeline, both the single-cycle unit and the floating point unit will need to write back result data simultaneously. Thus, the reservation station 20 must be inhibited from dispatching single-cycle micro
20 operations prior to when a floating point unit will want to write back its result data. The floating point units are therefore provided with a signal path to the reservation station 20 for signaling that certain types of operations should not be dispatched for one clock cycle. The logic in
25 the reservation station receives these signals and considers them in performing its scheduling algorithms. Because the reservation station of the preferred embodiment requires a

two clock-cycle pipeline for the scheduling of micro operations, the reservation station needs to receive the inhibit signal at least two clock cycles prior to the dispatch of any micro operations to be inhibited.

5 The first stage 41 of the floating point add functional unit 31, upon receiving a floating point add instruction, signals the reservation station not to issue any single-cycle micro operations. This prevents the reservation station from issuing a single-cycle micro operation to either
10 integer execution unit 34 or the wire unit 35 two cycles later which would then be attempting to write back result data at the same time as the floating point add result was being written back. Likewise, the second pipeline stage 49 of the multiply functional unit 32 signals the reservation
15 station 20 not to issue any single-cycle micro operations two cycles later which would result in the same writeback conflict with that unit.

One more floating point unit must be considered in the cluster receiving micro operations from Dispatch Port 0.
20 The floating point divide unit performs both floating point division and square root functions. This functional unit in the preferred embodiment is not pipelined. Floating point divides and square roots can have very long latencies, up to and including fifty to sixty clock cycles. While one
25 operation is being performed therein, no other micro operations may be dispatched to this functional unit 33.

Accordingly, the floating point divide unit upon receiving a micro operation signals the reservation station 20 to inhibit the dispatch of any subsequent micro operations to it until the floating point divide unit 33 has performed its result writeback. Alternatively, upon dispatching a divide, the reservation station logic may be designed to self-inhibit back-to-back divides. In addition, the floating point divide unit is capable of determining the number of clock cycles it will require for performing a received micro operation. With this information, the floating point divide unit 33 signals three cycles before writeback that no one cycle micro operations should be issued from the reservation station 20. Likewise, five, six and seven clock cycles prior to writeback, the floating point divide unit 33 signals that no three cycle, four or five cycle micro operations should be dispatched from the reservation station 20, respectively. It should be understood that when the reservation station is inhibited from dispatching only one class of micro operation (i.e., a three cycle floating point add), others may still be dispatched (i.e., single cycle operations or longer multiply operations).

The implementation of both the pipe lengthening of the floating point add unit and the inhibiting of the reservation station from issuing certain conflicting micro operations eliminates the potential problem of a writeback conflict by execution units sharing a common writeback path. The preferred embodiment of the present invention has been

described as being implemented in an out-of-order executing
microprocessor. The present invention may, however, be
implemented by any processor in which the above-described
writeback conflicts exist. Although the present invention
5 has been described in terms of preferred and implemented
embodiments, it will be appreciated that various
modifications and alterations might be made by those skilled
in the art without departing from the spirit and scope of the
present invention. The invention should, therefore, be
10 measured in terms of the claims which follow.



CLAIMS

1. An apparatus for extending the length of a pipeline of a first execution unit wherein said first execution unit shares a writeback path with a second execution unit, said first execution unit having at least a final pipeline stage, said second execution unit having a pipeline having at least one stage more than said first execution unit, said apparatus comprising:

buffering means coupled to said final pipeline stage of said first execution unit for delaying write back of a result of said first execution unit;

multiplexing means coupled to said buffering means and to said final pipeline stage of said first execution unit, said buffering means providing a first input to said multiplexing means and said final pipeline stage of said first execution unit providing a second input to said multiplexing means, said multiplexing means providing an output to said writeback path; and

writeback conflict detect means coupled to said first and second execution units for detecting if said execution units will encounter a writeback conflict, said writeback



conflict detect means being coupled to said multiplexing means for providing a control signal to said multiplexing means for selecting one of said first or second inputs to be the output provided to said writeback path.

2. The apparatus of claim 1 wherein said buffering means comprises a register coupled to receive the output of said final pipeline stage of said first execution unit.

3. The apparatus of claim 2 wherein said multiplexing means comprises multiplexer logic which receives as input the contents of said register and the output of said final pipeline stage of said first execution unit.

4. The apparatus of claim 3 wherein said writeback conflict detect means comprises comparing logic coupled to said first and second execution units for comparing the pipeline activity of said execution units to determine if said execution units will simultaneously write to said writeback path, said comparing logic providing a control signal to said multiplexer logic for selecting the input from said register to avoid a detected writeback conflict.



5. The apparatus of claim 1 wherein said first and second execution units are coupled to a reservation station, said reservation station having an n-cycle scheduling pipeline, said first execution unit having an m-cycle pipeline, said reservation station further being coupled to at least a third execution unit requiring only a single cycle to execute dispatched instructions, said third execution unit sharing said common writeback path with said first execution unit, said first execution unit further comprising means for signaling said reservation station to delay scheduling instructions for dispatch to said third execution unit to avoid a writeback conflict between said first and third execution units.

6. The apparatus of claim 1 wherein said first and second execution units are coupled to a reservation station, said reservation station having an n-cycle scheduling pipeline, said second execution unit having an m-cycle pipeline, said reservation station further being coupled to at least a third execution unit requiring only a single cycle to execute dispatched instructions, said third execution unit sharing said common writeback path with said second execution unit, said second execution unit further comprising means for signaling said reservation station to delay scheduling instructions for dispatch to said third execution unit to



- 27 -

avoid a writeback conflict between said second and third execution units.

7. The apparatus of claim 6 wherein said second execution unit signals said reservation station at the (m-n)th stage of the pipeline of said second execution unit.

8. The apparatus of claim 1 wherein said first execution unit has a pipeline length of N pipe stages and said second execution unit has a pipeline length of M pipe stages where M is greater than N, said buffering means comprising an (M-N) buffer depth.

9. The apparatus of claim 8 wherein (M-N) is greater than one (1).

10. The apparatus of claim 1 wherein said first execution unit comprises a floating point add unit having a three cycle pipeline and said second execution unit comprises an integer multiply unit having a four cycle pipeline.



- 28 -

11. An apparatus for extending the length of a pipeline of a first execution unit wherein said first execution unit shares a writeback path with a second execution unit, said first execution unit having at least a final pipeline stage, said second execution unit having a pipeline having at least one stage more than said first execution unit, said apparatus comprising:

a storage register coupled to said final pipeline stage of said first execution unit for latching the result of said first execution unit;

a multiplexer coupled to said storage register and to said final pipeline stage of said first execution unit, said storage register providing a first input to said multiplexer and said final pipeline stage of said first execution unit providing a second input to said multiplexer, said multiplexer providing an output to said writeback path; and

comparing logic coupled to said first and second execution units for detecting if said execution units will encounter a writeback conflict, said comparing logic being coupled to said multiplexer for providing a control signal to said multiplexer for selecting one of said first or second inputs to be the output provided to said writeback path.



- 29 -

12. The apparatus of claim 11 wherein said first and second execution units are coupled to a reservation station, said reservation station having an n-cycle scheduling pipeline, said second execution unit having an m-cycle pipeline, said reservation station further being coupled to at least a third execution unit requiring only a single cycle to execute dispatched instructions, said third execution unit sharing said common writeback path with said second execution unit, said second execution unit further comprising means for signaling said reservation station to delay scheduling instructions for dispatch to said third execution unit to avoid a writeback conflict between said second and third execution units.

13. The apparatus of claim 12 wherein said second execution unit signals said reservation station at the (m-n)th stage of the pipeline of said second execution unit.

14. The apparatus of claim 11 wherein said first execution unit comprises a floating point add unit having a three cycle pipeline and said second execution unit comprises an integer multiply unit having a four cycle pipeline.



- 30 -

15. For use in a microprocessor having at least first and second execution units which share a common writeback path, said first execution unit having at least a final pipeline stage and said second execution unit having a pipeline having at least one stage more than said pipeline of said first execution unit the method of avoiding writeback conflicts between said first and second execution units comprising the steps of:

providing said first execution unit with an additional buffer for delaying the output of said first execution, said additional buffer receiving the result from said final pipeline stage of said first execution unit;

comparing the pipeline activity of said first and second execution units to determine if both will attempt to use the common writeback path simultaneously; and

buffering the result from said first execution unit through said additional buffer when there would otherwise be a writeback conflict between said first and second execution units.

16. The method of claim 15 wherein said comparing and buffering steps comprise the steps of:

providing a multiplexer with the result from the final pipeline stage of said first execution unit;



- 31 -

simultaneously providing said multiplexer with the buffered result; and

selecting the appropriate signal from said multiplexer to write back over said common writeback path to avoid a writeback conflict.

17. The method of claim 15 wherein said first execution unit has a pipeline length of N pipe stages and said second execution unit has a pipeline length of M pipe stages where M is greater than N , said buffering step comprising buffering said result for $(M-N)$ pipe stages.

18. The method of claim 17 wherein $(M-N)$ is greater than one (1).

19. The method of claim 15 wherein said first and second execution units receive instructions from a reservation station having an n -cycle scheduling pipeline, said first execution unit having an m -cycle execution pipeline, said reservation station being further coupled to a third execution unit requiring only a single cycle to execute instructions, said method further comprising the step of said first execution unit, during the $(m-n)$ th stage of its



- 32 -

pipeline, signaling said reservation station not to dispatch any single cycle instructions.

20. The method of claim 19 wherein said microprocessor further comprises a fourth execution unit coupled to said reservation station and to said common writeback path, said fourth execution unit not being pipelined, the method comprising the steps of:

said fourth execution unit signaling said reservation station not to dispatch any instructions to said fourth execution unit until it has written a result over said write back path;

determining the number of clock cycles (X) that said fourth execution unit requires to carry out an instruction;

said fourth execution unit, during the (X-n)th cycle of execution signaling said reservation station not to dispatch any single cycle instructions; and

said fourth execution unit, during the (X-(m+n))th cycle of execution signaling said reservation station not to dispatch any m-cycle instructions.

21. For use in a microprocessor having at least first and second execution units which share a common writeback



- 33 -

path, said first execution unit having at least a final pipeline stage and said second execution unit having a pipeline having at least one stage more than said pipeline of said first execution unit, the method of avoiding writeback conflicts between said first and second execution units comprising the step of buffering the output of said first execution unit for an additional clock cycle when said first and second execution units would otherwise attempt to write back simultaneously.

22. A computer system comprising:

an input/output (I/O) means for providing a communications interface;

a memory means coupled to said I/O means for storing instructions and computer data;

data input means coupled to said I/O means for providing data input and data output to interface with a computer user; and

microprocessor means coupled to said I/O means for executing instructions and processing computer data, said microprocessor means further comprising a means for extending the length of a pipeline of a first execution unit wherein said first execution unit shares a writeback path with a second execution unit, said first execution unit having at



- 34 -

least a final pipeline stage, said second execution unit having a pipeline having at least one stage more than said first execution unit, said means for extending the length of a pipeline comprising:

buffering means coupled to said final pipeline stage of said first execution unit for delaying the result writeback from said first execution unit;

multiplexing means coupled to said buffering means and to said final pipeline stage of said first execution unit, said buffering means providing a first input to said multiplexing means and said final pipeline stage of said first execution unit providing a second input to said multiplexing means, said multiplexing means providing an output to said writeback path; and

writeback conflict detect means coupled to said first and second execution units for detecting if said execution units will encounter a writeback conflict, said writeback conflict detect means being coupled to said multiplexing means for providing a control signal to said multiplexing means for selecting one of said first or second inputs to be the output provided to said writeback path.

23. The computer system of claim 22 wherein said first and second execution units are coupled to a reservation




- 35 -

station, said reservation station having an n -cycle scheduling pipeline, said second execution unit having an m -cycle pipeline, said reservation station further being coupled to at least a third execution unit requiring only a single cycle to execute dispatched instructions, said third execution unit sharing said common writeback path with said second execution unit, said second execution unit further comprising means for signaling said reservation station to delay scheduling instructions for dispatch to said third execution unit to avoid a writeback conflict between said second and third execution units.

24. The computer system of claim 23 wherein said second execution unit signals said reservation station at the $(m-n)$ th stage of the pipeline of said second execution unit.

25. The computer system of claim 22 wherein said first execution unit has a pipeline length of N pipe stages and said second execution unit has a pipeline length of M pipe stages where M is greater than N , said buffering means comprising an $(M-N)$ buffer depth.

26. The computer system of claim 25 wherein $(M-N)$ is greater than one (1).



- 36 -

27. The computer system of claim 22 wherein said first execution unit comprises a floating point add unit having a three cycle pipeline and said second execution unit comprises an integer multiply unit having a four cycle pipeline.

28. A microprocessor comprising:

first and second pipelined execution units which may output their respective results during a same clock cycle or during different clock cycles; and

means, coupled to the first and second pipelined execution units, for delaying the first pipelined execution unit's result by at least one clock cycle when the first and second pipelined execution units' results would otherwise be output during the same clock cycle.

29. In a processor having first and second execution units coupled to receive instructions from, and write results back to, a common unit, the first execution unit having a final pipeline stage providing a first result during a given pipeline cycle, the second execution unit having at least one more pipeline stage than the first execution unit and including a final pipeline stage, the improvement comprising:



- 37 -

an extra pipeline stage coupled to receive the first result and to provide it as a delayed result during a subsequent pipeline cycle;

a multiplexer coupled to receive the first result and the delayed result, and coupled to the common unit; and

write back conflict detection logic coupled to the first and second execution units and to the multiplexer, to cause the multiplexer to select the delayed result for writeback to the common unit during the subsequent pipeline cycle if, in the given pipeline cycle, the final stages of both the first and second execution units will output results, and otherwise to write the first result back to the common unit during the given pipeline cycle.

30. An apparatus comprising:

an instruction issue unit;

a first execution unit, coupled to the instruction issue unit to execute instructions and to provide a first result during a first time period; and

a second execution unit coupled to the instruction issue unit to execute instructions and including,

a first pipeline stage coupled to the instruction issue unit to provide a second result, and.



- 38 -

a second pipeline stage coupled to the first pipeline stage to provide the second result as a delayed result during a second time period,

writeback conflict detection logic coupled to the first pipeline stage and to the first execution unit to detect whether the second result will be provided during the first time period, and

a multiplexer coupled to the first and second pipeline stages and to the detector, to write back the second result to the instruction issue unit if the second result will not be provided during the first time period, and to write back the delayed result to the instruction issue unit during the second time period otherwise.