

CS4617 Computer Architecture

Lecture10: Pipelining (continued)

Reference: Appendix C, Hennessy & Patterson

Dr J Vaughan

October 13, 2014

MIPS data path implementation (unpipelined)

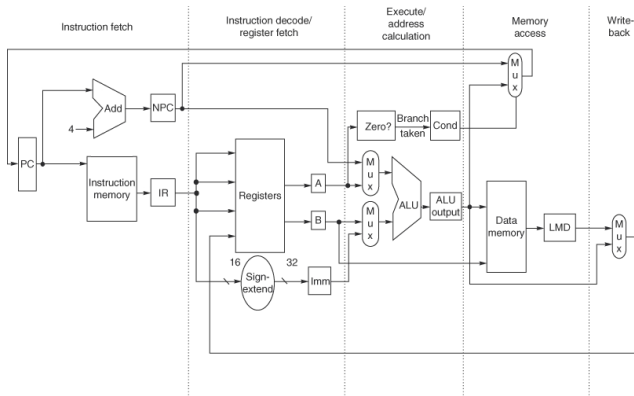


Figure C.21 The implementation of the MIPS data path allows every instruction to be executed in 4 or 5 clock cycles. Although the PC is shown in the portion of the data path that is used in instruction fetch and the registers are shown in the portion of the data path that is used in instruction decode/register fetch, both of these functional units are read as well as written by an instruction. Although we show these functional units in the cycle corresponding to where they are read, the PC is written during the memory access clock cycle and the registers are written during the write-back clock cycle. In both cases, the writes in later pipe stages are indicated by the multiplexer output (in memory access or write-back), which carries a value back to the PC or registers. These backward-flowing signals introduce much of the complexity of pipelining, since they indicate the possibility of hazards.

Basic Pipeline for MIPS

- ▶ Figure C21 is adapted by adding interstage registers to become the pipeline shown in Figure C22.
- ▶ The pipeline registers carry data and control from one stage to the next
- ▶ Values are copied along the registers until no longer needed
- ▶ The temporary registers used in the unpipelined processor are unsuitable as the values they contain could be overwritten before being completely used
- ▶ All registers needed to hold values temporarily between clock cycles within one instruction are contained in pipeline registers
- ▶ The pipeline registers carry data and control from one stage to the next
- ▶ A value needed in a later stage must be copied between pipeline registers until it is no longer needed

MIPS data path implementation (pipelined)

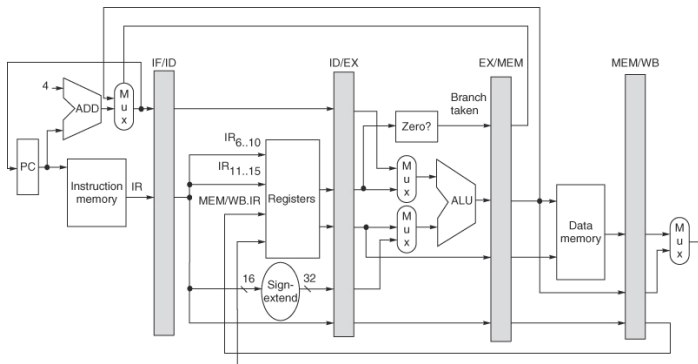


Figure C.22 The data path is pipelined by adding a set of registers, one between each pair of pipe stages. The registers serve to convey values and control information from one stage to the next. We can also think of the PC as a pipeline register, which sits before the IF stage of the pipeline, leading to one pipeline register for each pipe stage. Recall that the PC is an edge-triggered register written at the end of the clock cycle; hence, there is no race condition in writing the PC. The selection multiplexer for the PC has been moved so that the PC is written in exactly one stage (IF). If we didn't move it, there would be a conflict when a branch occurred, since two instructions would try to write different values into the PC. Most of the data paths flow from left to right, which is from earlier in time to later. The paths flowing from right to left (which carry the register write-back information and PC information on a branch) introduce complications into our pipeline.

- ▶ Using just one temporary register as in the unpipelined data path could cause values to be overwritten before all uses were completed
- ▶ For example, the field for a register operand for a write in a word or ALU operation comes from MEM/WB rather than IF/ID
- ▶ Any actions taken on behalf of an instruction occur between a pair of pipeline registers
- ▶ Figure C23 shows pipeline stage activities for various instruction types
- ▶ Actions in stages 1 and 2 are independent of instruction type as the instruction has not been decoded yet

MIPS pipeline IF & ID stage events

Stage	Any Instruction
IF	IF/ID.IR \leftarrow Mem[PC]; IF/ID.NPC, PC \leftarrow (if ((EX/MEM.opcode == branch) & EX/MEM.cond){EX/MEM.ALUOutput} else {PC+4});
ID	ID/EX.A \leftarrow Regs[IF/ID.IR[rs]]; ID/EX.B \leftarrow Regs[IF/ID.IR[rt]]; ID/EX.NPC \leftarrow IF/ID.NPC; ID/EX.IR \leftarrow IF/ID.IR; ID/EX.Imm \leftarrow sign-extend(IF/ID.IR[immediate field]);

Table: Figure C.23(a): Events on stages IF and ID of the MIPS pipeline

MIPS pipeline EX, MEM and WB events

Stage	ALU Instruction	Load or Store Instruction	Branch Instruction
EX	EX/MEM.IR \leftarrow ID/EX.IR; EX/MEM.ALUOutput \leftarrow ID/EX.A func ID/EX.B; or EX/MEM.ALUOutput \leftarrow ID/EX.A op ID/EX.Imm;	EX/MEM.IR to ID/EX.IR EX/MEM.ALUOutput \leftarrow ID/EX.A + ID/EX.Imm; EX/MEM.B \leftarrow ID/EX.B;	EX/MEM.ALUOutput \leftarrow ID/EX.NPC + (ID/EX.Imm $<<$ 2); EX/MEM.cond \leftarrow (ID/EX.A == 0);
MEM	MEM/WB.IR \leftarrow EX/MEM.IR; MEM/WB.ALUOutput \leftarrow EX/MEM.ALUOutput;	MEM/WB.IR \leftarrow EX/MEM.IR; MEM/WB.LMD \leftarrow Mem[EX/MEM.ALUOutput]; or Mem[EX/MEM.ALUOutput] \leftarrow EX/MEM.B;	
WB	Regs[MEM/WB.IR[rd]] \leftarrow MEM/WB.ALUOutput; Or Regs[MEM/WB.IR[rt]] \leftarrow MEM/WB.ALUOutput;	For load only: Regs [MEM/WB.IR[rt]] \leftarrow MEM/WB.LMD;	

Table: Figure C.23(b): Events on stages EX, MEM and WB of the MIPS pipeline

Figure C.23: Actions in the stages that are specific to the pipeline organization

- ▶ In IF, in addition to fetching the instruction and computing the new PC, we store the incremented PC both into the PC and into a pipeline register (NPC) for later use in computing the branch-target address
- ▶ This structure is the same as the organization in Figure C.22, where the PC is updated in IF from one of two sources
- ▶ In ID, we fetch the registers, extend the sign of the lower 16 bits of the IR (the immediate field), and pass along the IR and NPC
- ▶ During EX, we perform an ALU operation or an address calculation; we pass along the IR and the B register (if the instruction is a store)
- ▶ We also set the value of cond to 1 if the instruction is a taken branch
- ▶ During the MEM phase, we cycle the memory, write the PC if needed, and pass along values needed in the final pipe stage
- ▶ Finally, during WB, we update the register field from either the ALU output or the loaded value
- ▶ For simplicity, we always pass the entire IR from one stage to the next, although as an instruction proceeds down the pipeline, less and less of the IR is needed

MIPS pipeline control

- ▶ To control the simple pipeline, control the four multiplexers in the data path of Figure C22
- ▶ IF stage multiplexer controlled by EX/MEM.cond field
 - ▶ Chooses either PC+4 or EX/MEM.ALUOutput (the branch target) to write to the PC
- ▶ ALU stage multiplexers controlled by ID/EX.IR field
 - ▶ Top multiplexer set by whether or not instruction is a branch
 - ▶ Lower multiplexer set by whether or not instruction type is reg-reg ALU
- ▶ WB stage multiplexer controlled by whether instruction is Load or ALU operation
- ▶ Fifth multiplexer not shown in Figure C22
- ▶ Refer to Figure A22 for instruction formats

MIPS instruction formats


I-type instruction



Encodes: Loads and stores of bytes, half words, words, double words. All immediates ($rt \leftarrow rs \text{ op immediate}$)


Conditional branch instructions (rs is register, rd unused)
Jump register, jump and link register
($rd=0$, $rs=destination$, $immediate=0$)

R-type instruction



Register-register ALU operations: $rd \leftarrow rs \text{ funct } rt$
Function encodes the data path operation: Add, Sub, . . .
Read/write special registers and moves

J-type instruction



Jump and jump and link
Trap and return from exception

Figure A.22 Instruction layout for MIPS. All instructions are encoded in one of three types, with common fields in the same location in each format.

Write-back stage: the 5th multiplexer

- ▶ The destination field for WB is in a different place depending on instruction type
- ▶ Reg-reg ALU: $rd \leftarrow rx \text{ funct } rt$
Rd in bit positions 16-20 (counting from left)
- ▶ ALU immediate: $rt \leftarrow rs \text{ Op } immediate$
Rt in bit positions 11-15 (counting from left)
- ▶ Reg-reg ALU: $rt \leftarrow mem[rs + immediate]$
Rt in bit positions 11-15 (counting from left)
- ▶ The fifth multiplexer is needed to select either rd or rt as the specifier field for the register destination.

Implementing control for the MIPS pipeline

- ▶ *Instruction issue*: letting an instruction move from ID to EX of a pipeline
- ▶ An instruction that has passed from ID to EX is said to have *issued*
- ▶ In MIPS pipeline, all data hazards can be checked during ID
- ▶ If a hazard exists, the instruction is stalled before it is issued
- ▶ Any forwarding necessary can be determined during ID
- ▶ Detecting interlocks early in the pipeline reduces hardware complexity since the hardware never has to suspend an instruction that has updated the state of the processor unless the entire processor has stalled
- ▶ An alternative approach is to detect the hazard/forwarding at the beginning of a clock cycle that uses an operand (EX and MEM for this pipeline)

Example: Different approaches

1. Interlock for read after write (RAW) hazard
 - ▶ Source from load instruction (*load interlock*);
 - ▶ Check in ID
2. Forwarding paths to ALU inputs
 - ▶ Do during EX

Figure C24 shows different circumstances that must be handled

MIPS pipeline hazard detection comparisons

Situation	Example code sequence	Action
No dependence	LD R1,45(R2) DADD R5,R6,R7 DSUB R8,R6,R7 OR R9,R6,R7	No hazard possible because no dependence exists on R1 in the immediately following three instructions
Dependence requiring stall	LD R1,45(R2) DADD R5,R1,R7 DSUB R8,R6,R7 OR R9,R6,R7	Comparators detect the use of R1 in the DADD and stall the DADD (and DSUB and OR) before the DADD begins EX
Dependence overcome by forwarding	LD R1,45(R2) DADD R5,R6,R7 DSUB R8,R1,R7 OR R9,R6,R7	Comparators detect use of R1 in DSUB and forward result of load to ALU in time for DSUB to begin EX
Dependence with accesses in order	LD R1,45(R2) DADD R5,R6,R7 DSUB R8,R6,R7 OR R9,R1,R7	No action required because the read of R1 by OR occurs in the second half of the 1D phase, while the write of the loaded data occurred in the first half

Table: Figure C.24: Situations that the pipeline hazard detection hardware can see by comparing the destination and sources of adjacent instructions

Figure C.24: Legend

- ▶ Situations that the pipeline hazard detection hardware can see by comparing the destination and sources of adjacent instructions
- ▶ This table indicates that the only comparison needed is between the destination and the sources on the two instructions following the instruction that wrote the destination
- ▶ In the case of a stall, the pipeline dependences will look like the third case once execution continues
- ▶ Of course, hazards that involve R0 can be ignored since the register always contains 0, and the test above could be extended to do this

Load Interlock

- ▶ RAW hazard with source instruction = Load
- ▶ Load instruction is in EX when instruction that needs the data is in ID
- ▶ All possible hazards can be described in a small table that can translate directly to an implementation
- ▶ Figure C25 shows a table that detects all load interlocks when the instruction using the load result is in ID

MIPS pipeline hazard detection comparisons

Opcode field of ID/EX ($ID/EX.IR_{0..5}$)	Opcode field of IF/ID ($IF/ID.NR_{0..5}$)	Matching operand fields
Load	Register-register ALU	$ID/EX.IR[rt] == IF/ID.IR[rs]$
Load	Register-register ALU	$ID/EX.IR[rt] == IF/ID.IR[rt]$
Load	Load, store, ALU immediate, or branch	$ID/EX.IR[rt] == IF/ID.IR[rs]$

Table: Figure C.25: The logic to detect the need for load interlocks during the ID stage of an instruction requires three comparisons

Figure C.25: Legend

- ▶ The logic to detect the need for load interlocks during the ID stage of an instruction requires three comparisons
- ▶ Lines 1 and 2 of the table test whether the load destination register is one of the source registers for a register-register operation in ID
- ▶ Line 3 of the table determines if the load destination register is a source for a load or store effective address, an ALU immediate, or a branch test
- ▶ Remember that the IF/ID register holds the state of the instruction in ID, which potentially uses the load result, while ID/EX holds the state of the instruction in EX, which is the load instruction

After hazard has been detected

- ▶ Control unit must insert the stall and prevent instruction in IF and ID from advancing
- ▶ All control information is carried in the pipeline registers
- ▶ The instruction itself is carried and this is sufficient as all control is derived from it.
- ▶ Thus, when a hazard is detected
 1. Change the control portion of the ID/EX pipeline register to all zeros a NOP
 2. Recirculate the contents of the IF/ID registers to hold the stalled instruction
- ▶ In a pipeline with more complex hazards, apply the same ideas: detect the hazard by comparing some set of pipeline registers and shift in NOPs to prevent incorrect execution.

Forwarding Logic

- ▶ Similar to hazard treatment, but more cases to consider
- ▶ Pipeline register contain the data to be forwarded.
- ▶ Pipeline register contain source and destination register fields
- ▶ Forwarding is from ALU or data memory output to the ALU input, the data memory input, or the zero detection unit
- ▶ Can implement the forwarding by a comparison of the destination registers of the IR contained in the EX/MEM and MEM/WB registers.
- ▶ Figure C26 shows comparisons and possible forwarding when the destination of the forwarded result is an ALU input for the instruction currently in EX

MIPS pipeline forwarding comparisons (a)

Pipeline register containing source instruction	Opcode of source instruction	Pipeline register containing destination instruction	Opcode of destination instruction	Destination of forwarded result	Comparison (if equal then forward)
EX/MEM	Register-register ALU	ID/EX	Register-register ALU, ALU immediate, load, store, branch	Top ALU input	EX/MEM.IR[rd] == ID/EX.IR[rs]
EX/MEM	Register-register ALU	ID/EX	Register-register ALU	Bottom ALU input	EX/MEM.IR[rd] == ID/EX.IR[rt]
MEM/WB	Register-register ALU	ID/EX	Register-register ALU, ALU immediate, load, store, branch	Top ALU input	MEM/WB.IR[rd] == ID/EX.IR[rs]
MEM/WB	Register-register ALU	ID/EX	Register-register ALU	Bottom ALU input	MEM/WB.IR[rd] == ID/EX.IR[rt]
EX/MEM	ALU immediate	ID/EX	Register-register ALU, ALU immediate, load, store, branch	Top ALU input	EX/MEM.IR[rt] == ID/EX.IR[rs]

Table: Figure C.26(a): Forwarding of data to the two ALU inputs (for the instruction in EX)

MIPS pipeline forwarding comparisons(b)

Pipeline register containing source instruction	Opcode of source instruction	Pipeline register containing destination instruction	Opcode of destination instruction	Destination of forwarded result	Comparison (if equal then forward)
EX/MEM	ALU immediate	ID/EX	Register-register ALU	Bottom ALU input	EX/MEM.IR[rt] == ID/EX.IR[rt]
MEM/WB	ALU immediate	ID/EX	Register-register ALU, ALU immediate, load, store, branch	Top ALU input	MEM/WB.IR[rt] == ID/EX.IR[rs]
MEM/WB	ALU immediate	ID/EX	Register-register ALU	Bottom ALU input	MEM/WB.IR[rt] == ID/EX.IR[rt]
MEM/WB	Load	ID/EX	Register-register ALU, ALU immediate, load, store, branch	Top ALU input	MEM/WB.IR[rt] == ID/EX.IR[rs]
MEM/WB	Load	ID/EX	Register-register ALU	Bottom ALU input	MEM/WB.IR[rt] == ID/EX.IR[rt]

Table: Figure C.26(b): Forwarding of data to the two ALU inputs (for the instruction in EX)

Figure C.26: Legend

- ▶ Forwarding of data to the two ALU inputs (for the instruction in EX) can occur from the ALU result (in EX/MEM or in MEM/WB) or from the load result in MEM/WB
- ▶ There are 10 separate comparisons needed to tell whether a forwarding operation should occur
- ▶ The top and bottom ALU inputs refer to the inputs corresponding to the first and second ALU source operands, respectively, and are shown explicitly in Figure C.21 and in Figure C.27
- ▶ Remember that the pipeline latch for destination instruction in EX is ID/EX, while the source values come from the ALUOutput portion of EX/MEM or MEM/WB or the LMD portion of MEM/WB
- ▶ There is one complication not addressed by this logic: dealing with multiple instructions that write the same register
- ▶ For example, during the code sequence DADD R1, R2, R3; DADDI R1, R1, #2; DSUB R4, R3, R1, the logic must ensure that the DSUB instruction uses the result of the DADDI instruction rather than the result of the DADD instruction
- ▶ The logic shown above can be extended to handle this case by simply testing that forwarding from MEM/WB is enabled only when forwarding from EX/MEM is not enabled for the same input
- ▶ Because the DADDI result will be in EX/MEM, it will be forwarded, rather than the DADD result in MEM/WB

Forwarding needs

- ▶ Comparators and combinational logic to enable forwarding path
- ▶ Enlarged multiplexers at ALU inputs
- ▶ Connections from pipeline registers used to forward results
- ▶ Figure C.27 shows relevant segments of the pipelined data path.
- ▶ MIPS hazard detection and forwarding is relatively simple
- ▶ A floating-point extension is more complicated

MIPS result forwarding




Figure C.27 Forwarding of results to the ALU requires the addition of three extra inputs on each ALU multiplexer and the addition of three paths to the new inputs. The paths correspond to a bypass of: (1) the ALU output at the end of the EX, (2) the ALU output at the end of the MEM stage, and (3) the memory output at the end of the MEM stage.

Dealing with branches in the pipeline

- ▶ BEQ, BNE \implies test register for equality to another register, which may be R0
- ▶ Consider only BEQZ and BNEZ (zero test)
- ▶ Can complete decision by end of ID by moving the zero test into that cycle
- ▶ To take advantage of an early branch decision, must compute PC and NPC early
- ▶ Extra adder needed to calculate branch-target address during ID, because ALU is not usable until EX.
- ▶ Figure C.28 shows the revised pipeline data path
- ▶ Now only 1-clock cycle stall on branches
- ▶ However, an ALU instruction followed by a branch on its result will cause a data hazard stall
- ▶ Figure C.29 shows the branch part of the revised pipeline table from Figure C.23
- ▶ Some processors have more expensive branch hazards due to longer times required for calculation of the branch condition and the destination
- ▶ For example, this can occur if there are separate decode and register fetch stages
- ▶ The branch delay (length of control hazard) can become a significant branch penalty
- ▶ In general, the deeper the pipeline, the worse the branch penalty

Reducing the branch hazard stall




Figure C.28 The stall from branch hazards can be reduced by moving the zero test and branch-target calculation into the ID phase of the pipeline. Notice that we have made two important changes, each of which removes 1 cycle from the 3-cycle stall for branches. The first change is to move both the branch-target address calculation and the branch condition decision to the ID cycle. The second change is to write the PC of the instruction in the IF phase, using either the branch-target address computed during ID or the incremented PC computed during IF. In comparison, Figure C.22 obtained the branch-target address from the EX/MEM register and wrote the result during the MEM clock cycle. As mentioned in Figure C.22, the PC can be thought of as a pipeline register (e.g., as part of ID/IF), which is written with the address of the next instruction at the end of each IF cycle.

MIPS Revised pipeline structure

Pipe stage	Branch instruction
IF	$IF/ID.IR \leftarrow Mem[PC];$ $IF/ID.NPC, PC \leftarrow (if ((IF/ID.opcode == branch) \& (Regs[IF/ID.IR_6..10] op 0))$ $\{IF/ID.NPC + sign-extended (IF/ID.IR[immediate field] \ll 2)\} else \{PC + 4\});$
ID	$ID/EX.A \leftarrow Regs[IF/ID.IR_6..10]; ID/EX.B \leftarrow Regs[IF/ID.IR_{11..15}];$ $ID/EX.IR \leftarrow IF/ID.IR;$ $ID/EX.Imm \leftarrow (IF/ID.IR_{16})^{16} \# \# IF/ID.IR_{16..31}$
EX	
MEM	
WB	

Table: Figure C.29: Revised pipeline structure based on the original in Figure C.23

Figure C.29 Legend: revised pipeline structure is based on the original in Figure C.23

- ▶ It uses a separate adder, as in Figure C.28, to compute the branch-target address during ID
- ▶ The operations that are new or have changed are in bold
- ▶ Because the branch-target address addition happens during ID, it will happen for all instructions; the branch condition ($Regs[IF/ID.IR_{6..10}] \text{ op } 0$) will also be done for all instructions
- ▶ The selection of the sequential PC or the branch-target PC still occurs during IF, but it now uses values from the ID stage that correspond to the values set by the previous instruction
- ▶ This change reduces the branch penalty by 2 cycles:
 - ▶ one from evaluating the branch target and condition earlier
 - ▶ and one from controlling the PC selection on the same clock rather than on the next clock
- ▶ Since the value of cond is set to 0, unless the instruction in ID is a taken branch, the processor must decode the instruction before the end of ID
- ▶ Because the branch is done by the end of ID, the EX, MEM, and WB stages are unused for branches
- ▶ An additional complication arises for jumps that have a longer offset than branches
- ▶ We can resolve this by using an additional adder that sums the PC and lower 26 bits of the IR after shifting left by 2 bits