

ECG 700 Advanced Computer System Architecture Fall 2012

Lecture 3 – Instruction Level Parallelism and Its Exploitation

Mei Yang

Adapted from David Patterson's slides on graduate
computer architecture

Outline

- ▶ Introduction
- ▶ Basic Compiler Techniques for Exposing ILP
- ▶ Advanced Branch Prediction
- ▶ Dynamic Scheduling
- ▶ Hardware-Based Speculation
- ▶ Dynamic Scheduling, Multiple Issue, and Speculation
- ▶ Advanced Techniques for Instruction Delivery and Speculation
- ▶ Limitations of ILP
- ▶ Multithreading

2

Introduction

- ▶ Pipelining become universal technique in 1985
 - Overlaps execution of instructions
 - Exploits “Instruction Level Parallelism”
- ▶ Beyond this, there are two main approaches:
 - Hardware-based dynamic approaches
 - Used in server and desktop processors
 - Not used as extensively in PMD processors
 - Compiler-based static approaches
 - Not as successful outside of scientific applications

3

Instruction-Level Parallelism

- ▶ When exploiting instruction-level parallelism, goal is to maximize CPI
 - Pipeline CPI =
 - Ideal pipeline CPI +
 - Structural stalls +
 - Data hazard stalls +
 - Control stalls
- ▶ Parallelism with basic block is limited
 - Typical size of basic block = 3–6 instructions
 - Must optimize across branches

4

Data Dependence

- ▶ Loop-Level Parallelism
 - Unroll loop statically or dynamically
 - Use SIMD (vector processors and GPUs)

Example:

```
for (i=0; i<=999; i=i+1)
  x[i] = x[i] + y[i];
```

- ▶ Challenges:
 - Data dependency
 - Instruction j is data dependent on instruction i if
 - Instruction i produces a result that may be used by instruction j
 - Instruction j is data dependent on instruction k and instruction k is data dependent on instruction i
- ▶ Dependent instructions cannot be executed simultaneously

5

Data Dependence

- ▶ Dependencies are a property of programs
- ▶ Pipeline organization determines if dependence is detected and if it causes a stall
- ▶ Data dependence conveys:
 - Possibility of a hazard
 - Order in which results must be calculated
 - Upper bound on exploitable instruction level parallelism
- ▶ Dependencies that flow through memory locations are difficult to detect

6

Name Dependence

- ▶ Two instructions use the same name but no flow of information
 - Not a true data dependence, *but is a problem when reordering instructions*
 - *Antidependence*: instruction j writes a register or memory location that instruction i reads
 - Initial ordering (i before j) must be preserved
 - *Output dependence*: instruction i and instruction j write the same register or memory location
 - Ordering must be preserved
- ▶ To resolve, use renaming techniques

7

Other Factors

- ▶ Data Hazards
 - Read after write (RAW)
 - Write after write (WAW)
 - Write after read (WAR)
- ▶ Control Dependence
 - Ordering of instruction i with respect to a branch instruction
 - Instruction control dependent on a branch cannot be moved before the branch so that its execution is no longer controlled by the branch
 - An instruction not control dependent on a branch cannot be moved after the branch so that its execution is controlled by the branch

8

Examples

- Example 1:
 - DADDU R1,R2,R3
 - BEQZ R4,L
 - DSUBU R1,R1,R6
 - L: ...
 - OR R7,R1,R8
 - Example 2:
 - DADDU R1,R2,R3
 - BEQZ R12,skip
 - DSUBU R4,R5,R6
 - DADDU R5,R4,R9
 - skip:
 - OR R7,R8,R9
- ▶ OR instruction dependent on DADDU and DSUBU
 - ▶ Assume R4 isn't used after skip
 - Possible to move DSUBU before the branch

9

Compiler Techniques for Exposing ILP

- ▶ Pipeline scheduling
 - Separate dependent instruction from the source instruction by the pipeline latency of the source instruction
- ▶ Example:


```
for (i=999; i>=0; i=i-1)
  x[i] = x[i] + s;
```

Instruction producing result	Instruction using result	Latency in clock cycles
FP ALU op	Another FP ALU op	3
FP ALU op	Store double	2
Load double	FP ALU op	1
Load double	Store double	0

10

Pipeline Stalls

```

Loop:  L.D      F0,0(R1)
        stall
        ADD.D F4,F0,F2
        stall
        stall
        S.D F4,0(R1)
        DADDUI R1,R1,#-8
        stall (assume integer load latency is 1)
        BNE R1,R2,Loop
    
```

Instruction producing result	Instruction using result	Latency in clock cycles
FP ALU op	Another FP ALU op	3
FP ALU op	Store double	2
Load double	FP ALU op	1
Load double	Store double	0

11

Pipeline Scheduling

Scheduled code:

```

Loop:  L.D      F0,0(R1)
        DADDUI R1,R1,#-8
        ADD.D F4,F0,F2
        stall
        stall
        S.D F4,8(R1)
        BNE R1,R2,Loop
    
```

Instruction producing result	Instruction using result	Latency in clock cycles
FP ALU op	Another FP ALU op	3
FP ALU op	Store double	2
Load double	FP ALU op	1
Load double	Store double	0

12

Loop Unrolling

- ▶ Loop unrolling

- Unroll by a factor of 4 (assume # elements is divisible by 4)
- Eliminate unnecessary instructions

```
Loop:  L.D F0,0(R1)
      ADD.D F4,F0,F2
      S.D F4,0(R1) ;drop DADDUI & BNE
      L.D F6,-8(R1)
      ADD.D F8,F6,F2
      S.D F8,-8(R1) ;drop DADDUI & BNE
      L.D F10,-16(R1)
      ADD.D F12,F10,F2
      S.D F12,-16(R1) ;drop DADDUI & BNE
      L.D F14,-24(R1)
      ADD.D F16,F14,F2
      S.D F16,-24(R1)
      DADDUI R1,R1,#-32
      BNE R1,R2,Loop
```

- note: number of live registers vs. original loop

13

Loop Unrolling/Pipeline Scheduling

- ▶ Pipeline schedule the unrolled loop:

```
Loop:  L.D F0,0(R1)
      L.D F6,-8(R1)
      L.D F10,-16(R1)
      L.D F14,-24(R1)
      ADD.D F4,F0,F2
      ADD.D F8,F6,F2
      ADD.D F12,F10,F2
      ADD.D F16,F14,F2
      S.D F4,0(R1)
      S.D F8,-8(R1)
      DADDUI R1,R1,#-32
      S.D F12,16(R1)
      S.D F16,8(R1)
      BNE R1,R2,Loop
```

14

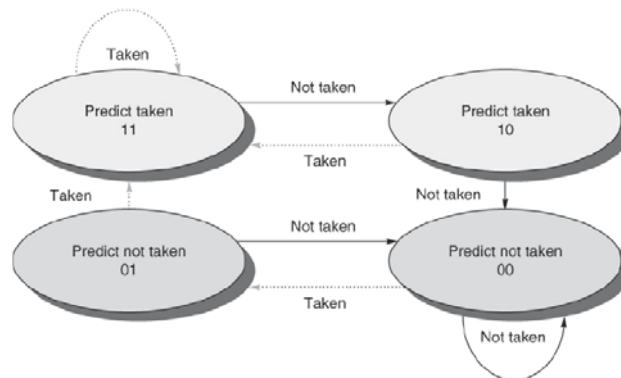
Strip Mining

- ▶ Unknown number of loop iterations?
 - Number of iterations = n
 - Goal: make k copies of the loop body
 - Generate pair of loops:
 - First executes $n \bmod k$ times
 - Second executes n / k times
 - “Strip mining”

15

Branch Prediction

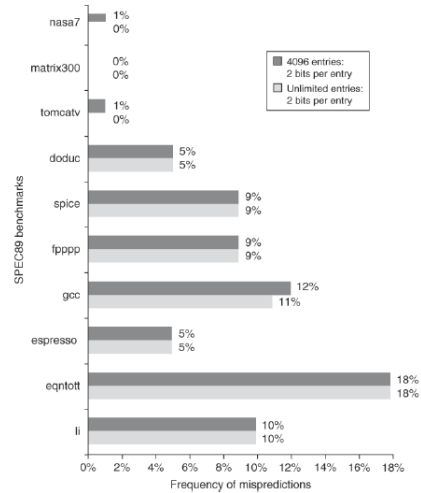
- ▶ Basic 2-bit predictor:
 - For each branch:
 - Predict taken or not taken
 - If the prediction is wrong two consecutive times, change prediction



16

Branch Prediction (cont'd)

- ▶ Branch history table (branch-prediction buffer) is used to store the branch prediction bit for each branch
 - Use low-order bits of branch PC to choose entry



17

Correlating Predictor

- ▶ Correlating predictor:
 - Multiple 2-bit predictors for each branch
 - One for each possible combination of outcomes of preceding n branches

Example:

```

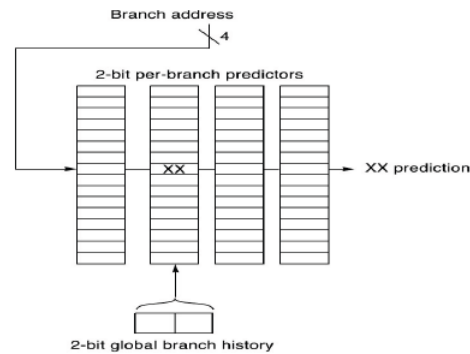
DADDIU R3, R1 #-2
BNEZ R3, L1 ; branch b1 (aa!=2)
If (aa==2)
  aa=0;
L1: DADDIU R3, R2, #-2
If (bb==2)
  bb=0;
L2: BNEZ R3, L2 ; branch b2 (bb!=2)
      DADD R2, R0, R0 ; bb=0
If (aa!=bb) {
L2: DSUBU R3, R1, R2 ; R3=aa-bb
      BEQZ R3, L3 ; branch b3 (aa==bb) which depends
                ; on the results of b1 and b2
    
```

18

Correlating Predictor (cont'd)

- (1,2) uses history of 1 branch and uses a 2-bit predictor
- (m, n) uses history of m branches to choose from 2^m branch predictors, each using n-bit

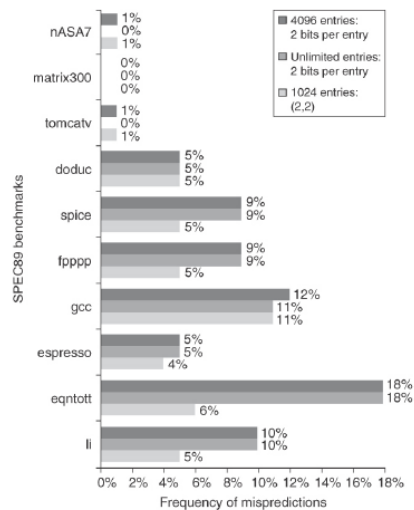
Example: (2, 2) uses 2 branches as history with 4 possibilities (T-T, NT-T, NT-NT, NT-T), each using a 2-bit predictor



19

Correlating Predictor (cont'd)

- ▶ Performance: With same number of state bits, (2,2) outperforms noncorrelating 2-bit predictor



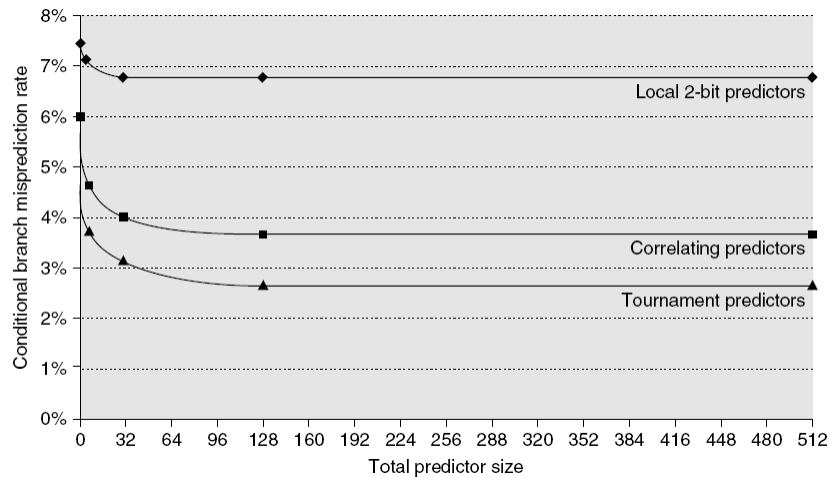
20

Tournament Predictor

- ▶ Local predictor:
 - Multiple 2-bit predictors for each branch
 - One for each possible combination of outcomes for the last n occurrences of this branch
- ▶ Tournament predictor:
 - Combine correlating predictor (global predictor) with local predictor
 - Choose one (local, global, or mix) which was most effective

21

Branch Prediction Performance



Branch predictor performance

22

Dynamic Scheduling

- ▶ Rearrange order of instructions to reduce stalls while maintaining data flow
- ▶ Dynamic scheduling implies:
 - Out-of-order execution
 - Out-of-order completion
- ▶ Advantages:
 - Compiler doesn't need to have knowledge of microarchitecture
 - Handles cases where dependencies are unknown at compile time
- ▶ Disadvantage:
 - Creates the possibility for WAR and WAW hazards
 - Substantial increase in hardware complexity
 - Complicates exceptions

23

Register Renaming

- ▶ Example:

```
DIV.D F0,F2,F4
ADD.D F6,F0,F8
S.D F6,0(R1)
SUB.D F8,F10,F14
MUL.D F6,F10,F8
```

antidependence

antidependence

+ name dependence with F6

24

Dynamic Scheduling

- ▶ To allow out-of-order execution, ID stage is split into two stages:
 - Issue: decode instructions, check for structural hazards
 - Read operands: wait until no data hazards, then read operands
- ▶ Scoreboarding allows out-of-order execution
- ▶ Tomasulo's Approach
 - Tracks when operands are available
 - Introduces register renaming in hardware
 - Minimizes WAW and WAR hazards

25

Register Renaming

- ▶ Register renaming is provided by reservation stations (RS)
 - Contains:
 - The instruction
 - Buffered operand values (when available)
 - Reservation station number of instruction providing the operand values
 - RS fetches and buffers an operand as soon as it becomes available (not necessarily involving register file)
 - Pending instructions designate the RS to which they will send their output
 - Result values broadcast on a result bus, called the common data bus (CDB)
 - Only the last output updates the register file
 - As instructions are issued, the register specifiers are renamed with the reservation station
 - May be more reservation stations than registers

26

Register Renaming

- ▶ Example:

DIV.D F0,F2,F4

ADD.D S,F0,F8

S.D S,0(R1)

SUB.D T,F10,F14

MUL.D F6,F10,T

- ▶ Now only RAW hazards remain, which can be strictly ordered

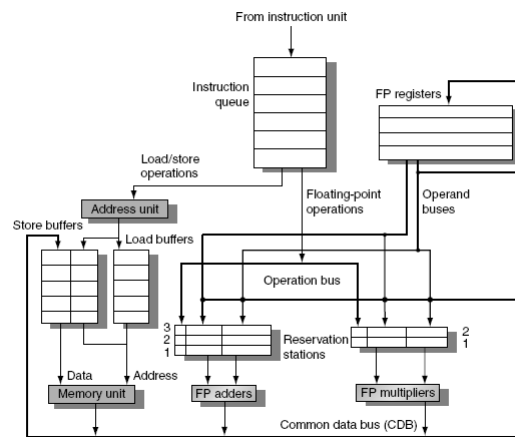
27

Tomasulo's Algorithm

- ▶ Load and store buffers

- Contain data and addresses, act like reservation stations

- ▶ Top-level design:



28

Tomasulo's Algorithm

- ▶ Three Steps:
 - Issue
 - Get next instruction from FIFO queue
 - If available RS, issue the instruction to the RS with operand values if available
 - If operand values not available, stall the instruction
 - Execute
 - When operand becomes available, store it in any reservation stations waiting for it
 - When all operands are ready, issue the instruction
 - Loads and store maintained in program order through effective address
 - No instruction allowed to initiate execution until all branches that proceed it in program order have completed
 - Write result
 - Write result on CDB into reservation stations and store buffers
 - (Stores must wait until address and value are received)

29

Reservation Station Components

- ▶ Op—Operation to perform in the unit (e.g., + or -)
- ▶ Qj, Qk—Reservation stations producing source registers
- ▶ Vj, Vk—Value of Source operands
- ▶ Rj, Rk—Flags indicating when Vj, Vk are ready
- ▶ Busy—Indicates reservation station is busy

- ▶ Register result status (Qi)—Indicates which reservation station (instruction) will write each register, if one exists. Blank when no pending instructions that will write that register.

30

Tomasulo Example

Instruction stream

Instruction status:

Instruction	<i>j</i>	<i>k</i>	Issue	Exec	Write
				Comp	Result
LD	F6	34+	R2		
LD	F2	45+	R3		
MULTD	F0	F2	F4		
SUBD	F8	F6	F2		
DIVD	F10	F0	F6		
ADDD	F6	F8	F2		

Load	Busy	Address
Load1	No	
Load2	No	
Load3	No	

3 Load/Buffers

Reservation Stations:

Time	Name	Busy	Op	<i>S1</i>	<i>S2</i>	<i>RS</i>	<i>RS</i>
				<i>Vj</i>	<i>Vk</i>	<i>Qj</i>	<i>Qk</i>
	Add1	No					
	Add2	No					
	Add3	No					
	Mult1	No					
	Mult2	No					

FU count down

3 FP Adder R.S.
2 FP Mult R.S.

Register result status:

Clock	F0	F2	F4	F6	F8	F10	F12	...	F30
0									

0
Clock cycle counter

Tomasulo Example Cycle 1

Instruction status:

Instruction	<i>j</i>	<i>k</i>	Issue	Exec	Write
				Comp	Result
LD	F6	34+	R2	1	
LD	F2	45+	R3		
MULTD	F0	F2	F4		
SUBD	F8	F6	F2		
DIVD	F10	F0	F6		
ADDD	F6	F8	F2		

Load	Busy	Address
Load1	Yes	34+R2
Load2	No	
Load3	No	

Reservation Stations:

Time	Name	Busy	Op	<i>S1</i>	<i>S2</i>	<i>RS</i>	<i>RS</i>
				<i>Vj</i>	<i>Vk</i>	<i>Qj</i>	<i>Qk</i>
	Add1	No					
	Add2	No					
	Add3	No					
	Mult1	No					
	Mult2	No					

Register result status:

Clock	F0	F2	F4	F6	F8	F10	F12	...	F30
1				Load1					

Tomasulo Example Cycle 2

Instruction status:

Instruction	j	k	Exec Write		Busy	Address
			Issue	Comp Result		
LD	F6	34+	R2	1	Yes	34+R2
LD	F2	45+	R3	2	Yes	45+R3
MULTD	F0	F2	F4		No	
SUBD	F8	F6	F2		No	
DIVD	F10	F0	F6		No	
ADDD	F6	F8	F2		No	

Reservation Stations:

Time	Name	Busy	Op	S1 Vj	S2 Vk	RS Qj	RS Qk
Add1		No					
Add2		No					
Add3		No					
Mult1		No					
Mult2		No					

Register result status:

Clock	F0	F2	F4	F6	F8	F10	F12	...	F30
2	FU	Load2							Load1

Note: Can have multiple loads outstanding

Tomasulo Example Cycle 3

Instruction status:

Instruction	j	k	Exec Write		Busy	Address
			Issue	Comp Result		
LD	F6	34+	R2	1	Yes	34+R2
LD	F2	45+	R3	2	Yes	45+R3
MULTD	F0	F2	F4	3	No	
SUBD	F8	F6	F2		No	
DIVD	F10	F0	F6		No	
ADDD	F6	F8	F2		No	

Reservation Stations:

Time	Name	Busy	Op	S1 Vj	S2 Vk	RS Qj	RS Qk
Add1		No					
Add2		No					
Add3		No					
Mult1		Yes	MULTD		R(F4)	Load2	
Mult2		No					

Register result status:

Clock	F0	F2	F4	F6	F8	F10	F12	...	F30
3	Mult1	Load2							Load1

- Note: registers names are removed ("renamed") in Reservation Stations; MULT issued
- Load1 completing; what is waiting for Load1?

Tomasulo Example Cycle 4

Instruction status:

Instruction	j	k	Issue	Exec		Write	Busy	Address
				Comp	Result			
LD	F6	34+	R2	1	3	4	No	
LD	F2	45+	R3	2	4		Yes	45+R3
MULTD	F0	F2	F4	3			No	
SUBD	F8	F6	F2	4			No	
DIVD	F10	F0	F6				No	
ADDD	F6	F8	F2				No	

Reservation Stations:

Time	Name	Busy	Op	S1 Vj	S2 Vk	RS Qj	RS Qk
Add1	Yes	SUBD	M(A1)				Load2
Add2	No						
Add3	No						
Mult1	Yes	MULTD			R(F4)	Load2	
Mult2	No						

Register result status:

Clock	F0	F2	F4	F6	F8	F10	F12	...	F30
4	FU	Mult1	Load2		M(A1)	Add1			

- Load2 completing; what is waiting for Load2?

Tomasulo Example Cycle 5

Instruction status:

Instruction	j	k	Issue	Exec		Write	Busy	Address
				Comp	Result			
LD	F6	34+	R2	1	3	4	No	
LD	F2	45+	R3	2	4	5	No	
MULTD	F0	F2	F4	3			No	
SUBD	F8	F6	F2	4			No	
DIVD	F10	F0	F6	5			No	
ADDD	F6	F8	F2				No	

Reservation Stations:

Time	Name	Busy	Op	S1 Vj	S2 Vk	RS Qj	RS Qk
3	Add1	Yes	SUBD	M(A1)	M(A2)		
Add2	No						
Add3	No						
7	Mult1	Yes	MULTD	M(A2)	R(F4)		
Mult2	Yes	DIVD		M(A1)	Mult1		

Register result status:

Clock	F0	F2	F4	F6	F8	F10	F12	...	F30
5	FU	Mult1	M(A2)		M(A1)	Add1	Mult2		

- Timer starts down for Add1, Mult1

Tomasulo Example Cycle 6

Instruction status:

Instruction	j	k	Issue	Exec Write		Busy	Address
				Comp	Result		
LD	F6	34+	R2	1	3	4	Load1
LD	F2	45+	R3	2	4	5	Load2
MULTD	F0	F2	F4	3			Load3
SUBD	F8	F6	F2	4			
DIVD	F10	F0	F6	5			
ADDD	F6	F8	F2	6			

Reservation Stations:

Time	Name	Busy	Op	S1 Vj	S2 Vk	RS Qj	RS Qk
2	Add1	Yes	SUBD	M(A1)	M(A2)		
	Add2	Yes	ADDD		M(A2)	Add1	
	Add3	No					
6	Mult1	Yes	MULTD	M(A2)	R(F4)		
	Mult2	Yes	DIVD		M(A1)	Mult1	

Register result status:

Clock	F0	F2	F4	F6	F8	F10	F12	...	F30
6	FU	Mult1	M(A2)		Add2	Add1	Mult2		

- Issue ADDD here despite name dependency on F6?

Tomasulo Example Cycle 7

Instruction status:

Instruction	j	k	Issue	Exec Write		Busy	Address
				Comp	Result		
LD	F6	34+	R2	1	3	4	Load1
LD	F2	45+	R3	2	4	5	Load2
MULTD	F0	F2	F4	3			Load3
SUBD	F8	F6	F2	4			
DIVD	F10	F0	F6	5			
ADDD	F6	F8	F2	6			

Reservation Stations:

Time	Name	Busy	Op	S1 Vj	S2 Vk	RS Qj	RS Qk
1	Add1	Yes	SUBD	M(A1)	M(A2)		
	Add2	Yes	ADDD		M(A2)	Add1	
	Add3	No					
5	Mult1	Yes	MULTD	M(A2)	R(F4)		
	Mult2	Yes	DIVD		M(A1)	Mult1	

Register result status:

Clock	F0	F2	F4	F6	F8	F10	F12	...	F30
7	FU	Mult1	M(A2)		Add2	Add1	Mult2		

Tomasulo Example Cycle 8

Instruction status:

Instruction	j	k	Issue	Exec Write		Load1	Load2	Load3	Busy	Address
				Comp	Result					
LD	F6	34+	R2	1	3	4			No	
LD	F2	45+	R3	2	4	5			No	
MULTD	F0	F2	F4	3					No	
SUBD	F8	F6	F2	4	8					
DIVD	F10	F0	F6	5						
ADDD	F6	F8	F2	6						

Reservation Stations:

Time	Name	Busy	Op	S1 Vj	S2 Vk	RS Qj	RS Qk
0	Add1	Yes	SUBD	M(A1)	M(A2)		
	Add2	Yes	ADDD		M(A2)	Add1	
	Add3	No					
4	Mult1	Yes	MULTD	M(A2)	R(F4)		
	Mult2	Yes	DIVD		M(A1)	Mult1	

Register result status:

Clock	F0	F2	F4	F6	F8	F10	F12	...	F30
8	FU	Mult1	M(A2)		Add2	Add1	Mult2		

- Add1 (SUBD) completing; what is waiting for it?

Tomasulo Example Cycle 9

Instruction status:

Instruction	j	k	Issue	Exec Write		Load1	Load2	Load3	Busy	Address
				Comp	Result					
LD	F6	34+	R2	1	3	4			No	
LD	F2	45+	R3	2	4	5			No	
MULTD	F0	F2	F4	3					No	
SUBD	F8	F6	F2	4	8	9				
DIVD	F10	F0	F6	5						
ADDD	F6	F8	F2	6						

Reservation Stations:

Time	Name	Busy	Op	S1 Vj	S2 Vk	RS Qj	RS Qk
	Add1	No					
3	Add2	Yes	ADDD	(M-M)	M(A2)		
	Add3	No					
3	Mult1	Yes	MULTD	M(A2)	R(F4)		
	Mult2	Yes	DIVD		M(A1)	Mult1	

Register result status:

Clock	F0	F2	F4	F6	F8	F10	F12	...	F30
9	FU	Mult1	M(A2)		Add2	(M-M)	Mult2		

Tomasulo Example Cycle 10

Instruction status:

Instruction	j	k	Exec Write			Load1	Load2	Load3	Busy	Address
			Issue	Comp	Result					
LD	F6	34+	R2	1	3	4		No		
LD	F2	45+	R3	2	4	5		No		
MULTD	F0	F2	F4	3				No		
SUBD	F8	F6	F2	4	8	9				
DIVD	F10	F0	F6	5						
ADDD	F6	F8	F2	6						

Reservation Stations:

Time	Name	Busy	Op	S1	S2	RS	RS
				Vj	Vk	Qj	Qk
Add1		No					
2 Add2	Yes	ADDD	(M-M)	M(A2)			
Add3	No						
2 Mult1	Yes	MULTD	M(A2)	R(F4)			
Mult2	Yes	DIVD		M(A1)	Mult1		

Register result status:

Clock	F0	F2	F4	F6	F8	F10	F12	...	F30
10	Mult1	M(A2)		Add2	(M-M)	Mult2			

Tomasulo Example Cycle 11

Instruction status:

Instruction	j	k	Exec Write			Load1	Load2	Load3	Busy	Address
			Issue	Comp	Result					
LD	F6	34+	R2	1	3	4		No		
LD	F2	45+	R3	2	4	5		No		
MULTD	F0	F2	F4	3				No		
SUBD	F8	F6	F2	4	8	9				
DIVD	F10	F0	F6	5						
ADDD	F6	F8	F2	6						

Reservation Stations:

Time	Name	Busy	Op	S1	S2	RS	RS
				Vj	Vk	Qj	Qk
Add1		No					
1 Add2	Yes	ADDD	(M-M)	M(A2)			
Add3	No						
1 Mult1	Yes	MULTD	M(A2)	R(F4)			
Mult2	Yes	DIVD		M(A1)	Mult1		

Register result status:

Clock	F0	F2	F4	F6	F8	F10	F12	...	F30
11	Mult1	M(A2)		Add2	(M-M)	Mult2			

Tomasulo Example Cycle 12

Instruction status:

Instruction	j	k	Exec Write			Busy	Address	
			Issue	Comp	Result			
LD	F6	34+	R2	1	3	4	Load1	No
LD	F2	45+	R3	2	4	5	Load2	No
MULTD	F0	F2	F4	3	12		Load3	No
SUBD	F8	F6	F2	4	8	9		
DIVD	F10	F0	F6	5				
ADDD	F6	F8	F2	6	12			

Reservation Stations:

Time	Name	Busy	Op	S1	S2	RS	RS
				Vj	Vk	Qj	Qk
	Add1	No					
0	Add2	Yes	ADDD	M(A2)	M(A2)		
	Add3	No					
0	Mult1	Yes	MULTD	M(A2)	R(F4)		
	Mult2	Yes	DIVD	M(A1)	Mult1		

Register result status:

Clock	F0	F2	F4	F6	F8	F10	F12	...	F30
12	FU	Mult1	M(A2)		Add2	(M-M)	Mult2		

- Add2 (ADDD) completing; what is waiting for it?
- Mult1 (MULTD) completing; what is waiting for it?

Tomasulo Example Cycle 13

Instruction status:

Instruction	j	k	Exec Write			Busy	Address	
			Issue	Comp	Result			
LD	F6	34+	R2	1	3	4	Load1	No
LD	F2	45+	R3	2	4	5	Load2	No
MULTD	F0	F2	F4	3	12	13	Load3	No
SUBD	F8	F6	F2	4	8	9		
DIVD	F10	F0	F6	5				
ADDD	F6	F8	F2	6	12	13		

Reservation Stations:

Time	Name	Busy	Op	S1	S2	RS	RS
				Vj	Vk	Qj	Qk
	Add1	No					
	Add2	No					
	Add3	No					
	Mult1	No					
13	Mult2	Yes	DIVD	M*F4	M(A1)		

Register result status:

Clock	F0	F2	F4	F6	F8	F10	F12	...	F30
13	FU	M*F4	M(A2)		(M-M+M)	(M-M)	Mult2		

- Write results of ADDD and MULT here?

Tomasulo Example Cycle 25

LD	F2	45+	R3	2	4	5	Load2	No
MULTD	F0	F2	F4	3	12	13	Load3	No
SUBD	F8	F6	F2	4	8	9		
DIVD	F10	F0	F6	5				
ADDD	F6	F8	F2	6	12	13		

Reservation Stations:

Time	Name	Busy	Op	S1 Vj	S2 Vk	RS Qj	RS Qk
	Add1	No					
	Add2	No					
	Add3	No					
	Mult1	No					
1	Mult2	Yes	DIVD	M*F4	M(A1)		

Register result status:

Clock	F0	F2	F4	F6	F8	F10	F12	...	F30
25	M*F4	M(A2)		(M-M+N)	(M-M)	Mult2			

Tomasulo Example Cycle 26

Instruction status:

Instruction	j	k	Exec			Busy	Address	
			Issue	Comp	Result			
LD	F6	34+	R2	1	3	4	Load1	No
LD	F2	45+	R3	2	4	5	Load2	No
MULTD	F0	F2	F4	3	12	13	Load3	No
SUBD	F8	F6	F2	4	8	9		
DIVD	F10	F0	F6	5	26			
ADDD	F6	F8	F2	6	12	13		

Reservation Stations:

Time	Name	Busy	Op	S1 Vj	S2 Vk	RS Qj	RS Qk
	Add1	No					
	Add2	No					
	Add3	No					
	Mult1	No					
0	Mult2	Yes	DIVD	M*F4	M(A1)		

Register result status:

Clock	F0	F2	F4	F6	F8	F10	F12	...	F30
26	M*F4	M(A2)		(M-M+N)	(M-M)	Mult2			

- Mult2 (DIVD) is completing; what is waiting for it?

Tomasulo Example Cycle 27

Instruction status:

Instruction	j	k	Exec Write			Busy	Address
			Issue	Comp	Result		
LD	F6	34+	R2	1	3	4	No
LD	F2	45+	R3	2	4	5	No
MULTD	F0	F2	F4	3	12	13	No
SUBD	F8	F6	F2	4	8	9	
DIVD	F10	F0	F6	5	26	27	
ADD	F6	F8	F2	6	12	13	

Reservation Stations:

Time	Name	Busy	Op	S1		S2		RS	
				Vj	Vk	Qj	Qk		
Add1		No							
Add2		No							
Add3		No							
Mult1		No							
Mult2	Yes		DIVD	M*F4	M(A1)				

Register result status:

Clock	F0	F2	F4	F6	F8	F10	F12	...	F30
27	FU	M*F4	M(A2)		(M-M+N	(M-M)	Result		

- Once again: In-order issue, out-of-order execution and out-of-order completion.

Tomasulo Algorithm

Instruction state	Wait until	Action or bookkeeping
Issue FP operation	Station r empty	if (RegisterStat[rs].Qi != 0) {RS[r].Qj ← RegisterStat[rs].Qi} else {RS[r].Vj ← Regs[rs]; RS[r].Qj ← 0}; if (RegisterStat[rt].Qi != 0) {RS[r].Qk ← RegisterStat[rt].Qi} else {RS[r].Vk ← Regs[rt]; RS[r].Qk ← 0}; RS[r].Busy ← yes; RegisterStat[rd].Q ← r;
Load or store	Buffer r empty	if (RegisterStat[rs].Qi != 0) {RS[r].Qj ← RegisterStat[rs].Qi} else {RS[r].Vj ← Regs[rs]; RS[r].Qj ← 0}; RS[r].A ← imm; RS[r].Busy ← yes;
Load only		RegisterStat[rt].Qi ← r;
Store only		if (RegisterStat[rt].Qi != 0) {RS[r].Qk ← RegisterStat[rs].Qi} else {RS[r].Vk ← Regs[rt]; RS[r].Qk ← 0};
Execute FP operation	(RS[r].Qj = 0) and (RS[r].Qk = 0)	Compute result: operands are in Vj and Vk
Load/store step 1	RS[r].Qj = 0 & r is head of load-store queue	RS[r].A ← RS[r].Vj + RS[r].A;
Load step 2	Load step 1 complete	Read from Mem[RS[r].A]
Write result FP operation or load	Execution complete at r & CDB available	∀x (if (RegisterStat[x].Qi = r) {Regs[x] ← result; RegisterStat[x].Qi ← 0}); ∀x (if (RS[x].Qj = r) {RS[x].Vj ← result; RS[x].Qj ← 0}); ∀x (if (RS[x].Qk = r) {RS[x].Vk ← result; RS[x].Qk ← 0}); RS[r].Busy ← no;
Store	Execution complete at r & RS[r].Qk = 0	Mem[RS[r].A] ← RS[r].Vk; RS[r].Busy ← no;

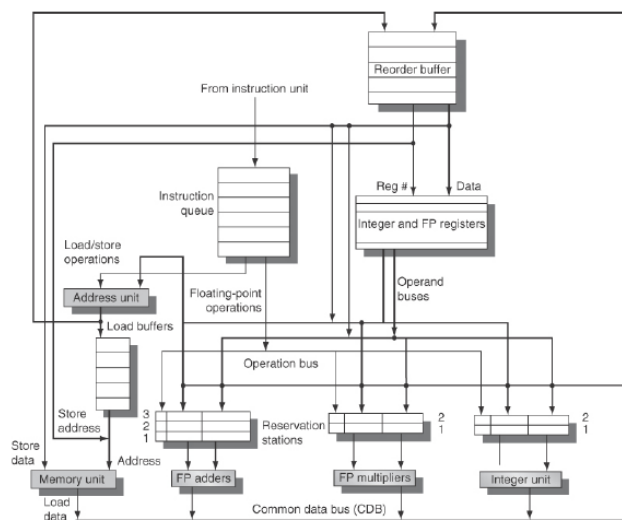
48

Hardware-Based Speculation

- ▶ Execute instructions along predicted execution paths but only commit the results if prediction was correct
- ▶ Instruction commit: allowing an instruction to update the register file when instruction is no longer speculative
- ▶ Need an additional piece of hardware to prevent any irrevocable action until an instruction commits
 - I.e. updating state or taking an execution

49

Hardware-Based Speculation



50

Reorder Buffer

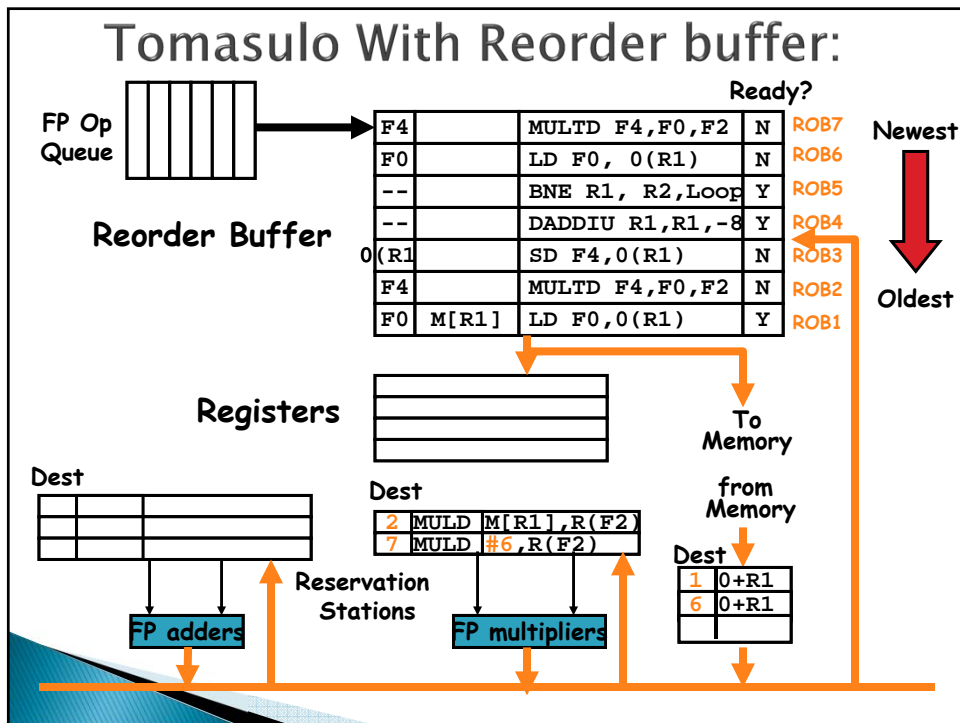
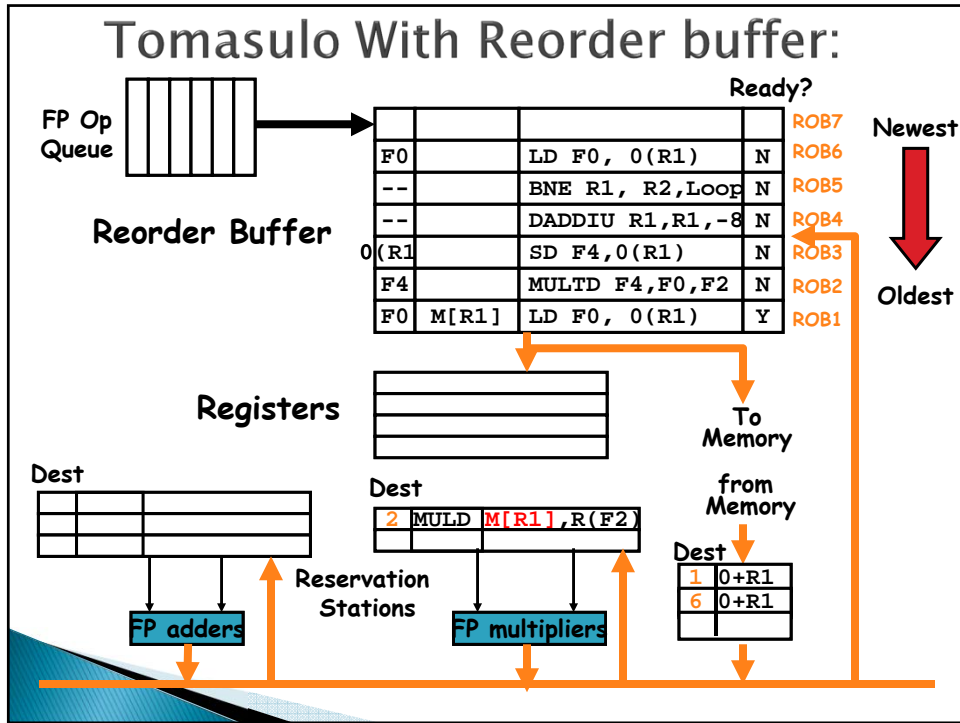
- ▶ Reorder buffer – holds the result of instruction between completion and commit
- ▶ Four fields:
 - Instruction type: branch/store/register
 - Destination field: register number
 - Value field: output value
 - Ready field: completed execution?
- ▶ Modify reservation stations:
 - Operand source is now reorder buffer instead of functional unit

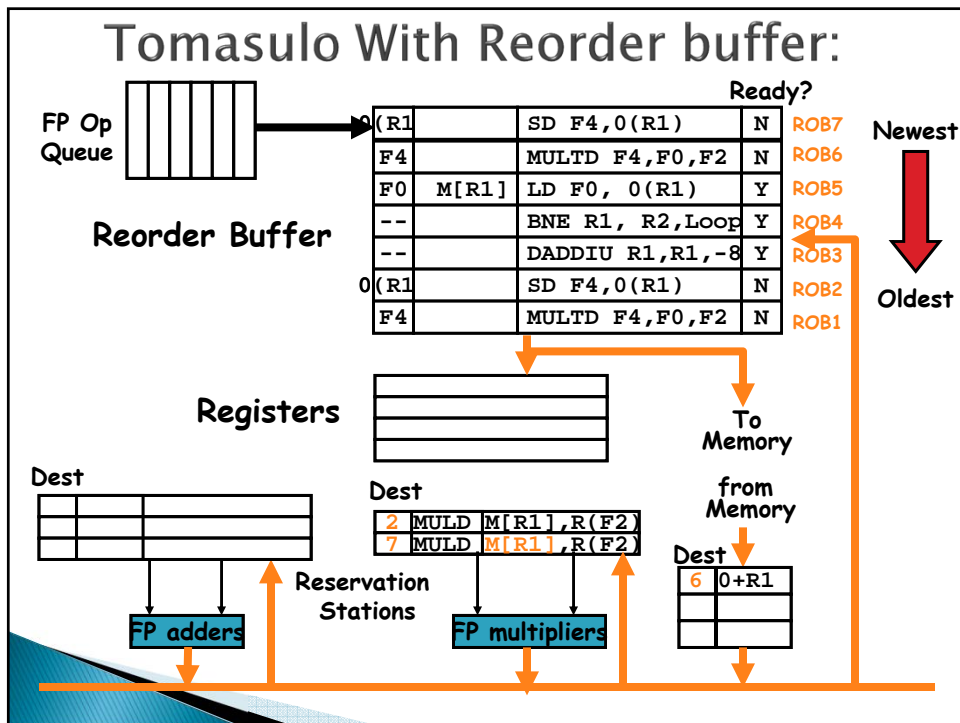
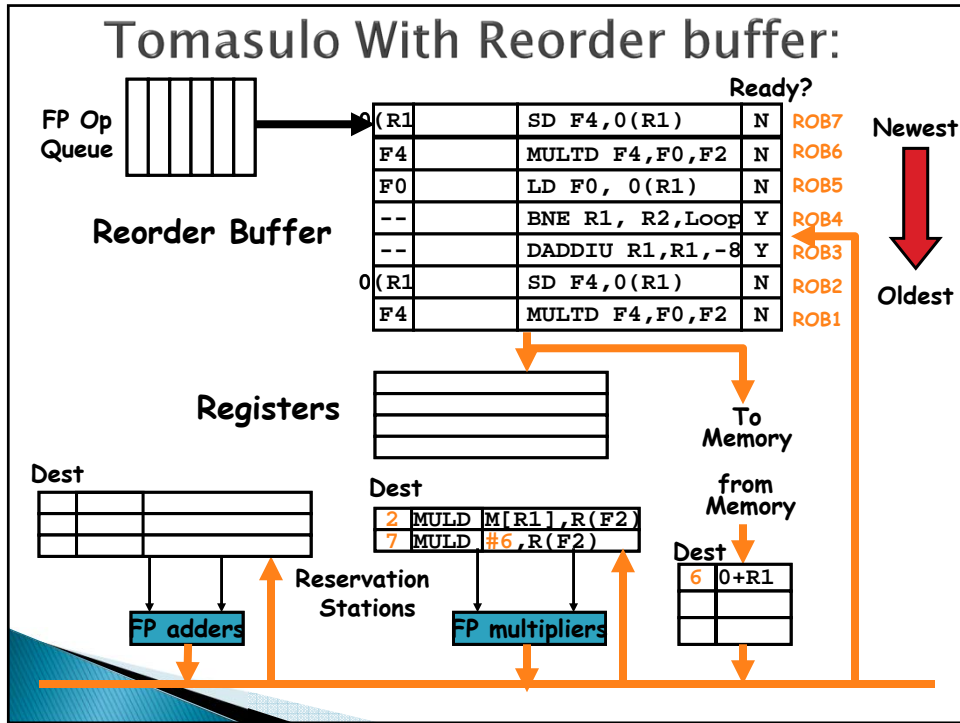
51

Reorder Buffer

- ▶ Register values and memory values are not written until an instruction commits
- ▶ On misprediction:
 - Speculated entries in ROB are cleared
- ▶ Exceptions:
 - Not recognized until it is ready to commit

52





Avoiding Memory Hazards

- ▶ WAW and WAR hazards through memory are eliminated with speculation because actual updating of memory occurs in order, when a store is at head of the ROB, and hence, no earlier loads or stores can still be pending
- ▶ RAW hazards through memory are maintained by two restrictions:
 1. not allowing a load to initiate the second step of its execution if any active ROB entry occupied by a store has a Destination field that matches the value of the A field of the load, and
 2. maintaining the program order for the computation of an effective address of a load with respect to all earlier stores.
- ▶ these restrictions ensure that any load that accesses a memory location written to by an earlier store cannot perform the memory access until the store has written the data

Multiple Issue and Static Scheduling

- ▶ To achieve $CPI < 1$, need to complete multiple instructions per clock
- ▶ Solutions:
 - Statically scheduled superscalar processors
 - VLIW (very long instruction word) processors
 - dynamically scheduled superscalar processors

Multiple Issue

Common name	Issue structure	Hazard detection	Scheduling	Distinguishing characteristic	Examples
Superscalar (static)	Dynamic	Hardware	Static	In-order execution	Mostly in the embedded space: MIPS and ARM, including the ARM Cortex A8
Superscalar (dynamic)	Dynamic	Hardware	Dynamic	Some out-of-order execution, but no speculation	None at the present
Superscalar (speculative)	Dynamic	Hardware	Dynamic with speculation	Out-of-order execution with speculation	Intel Core i3, i5, i7; AMD Phenom; IBM Power 7
VLIW/LIW	Static	Primarily software	Static	All hazards determined and indicated by compiler (often implicitly)	Most examples are in signal processing, such as the TI C6x
EPIC	Primarily static	Primarily software	Mostly static	All hazards determined and indicated explicitly by the compiler	Itanium

63

VLIW Processors

- ▶ Package multiple operations into one instruction
- ▶ Example VLIW processor:
 - One integer instruction (or branch)
 - Two independent floating-point operations
 - Two independent memory references
- ▶ Must be enough parallelism in code to fill the available slots

64

Recall: Unrolled Loop that Minimizes Stalls for Scalar

```

1 Loop: L.D   F0,0(R1)
2       L.D   F6,-8(R1)
3       L.D   F10,-16(R1)
4       L.D   F14,-24(R1)
5       ADD.D F4,F0,F2
6       ADD.D F8,F6,F2
7       ADD.D F12,F10,F2
8       ADD.D F16,F14,F2
9       S.D   0(R1),F4
10      S.D   -8(R1),F8
11      S.D   -16(R1),F12
12      DSUBUI R1,R1,#32
13      BNEZ  R1,LOOP
14      S.D   8(R1),F16 ; 8-32 = -24
    
```

L.D to ADD.D: 1 Cycle
ADD.D to S.D: 2 Cycles

14 clock cycles, or 3.5 per iteration

Loop Unrolling in VLIW

Memory reference 1	Memory reference 2	FP operation 1	FP op. 2	Int. op/branch	Clock
L.D F0,0(R1)	L.D F6,-8(R1)				1
L.D F10,-16(R1)	L.D F14,-24(R1)				2
L.D F18,-32(R1)	L.D F22,-40(R1)	ADD.D F4,F0,F2	ADD.D F8,F6,F2		3
L.D F26,-48(R1)		ADD.D F12,F10,F2	ADD.D F16,F14,F2		4
		ADD.D F20,F18,F2	ADD.D F24,F22,F2		5
S.D 0(R1),F4	S.D -8(R1),F8	ADD.D F28,F26,F2			6
S.D -16(R1),F12	S.D -24(R1),F16				7
S.D -32(R1),F20	S.D -40(R1),F24			DSUBUI R1,R1,#48	8
S.D -0(R1),F28				BNEZ R1,LOOP	9

Unrolled 7 times to avoid delays

7 results in 9 clocks, or 1.3 clocks per iteration (1.8X)

Average: 2.5 ops per clock, 50% efficiency

Note: Need more registers in VLIW (15 vs. 6 in SS)

VLIW Processors

- ▶ Disadvantages:
 - Statically finding parallelism
 - Code size
 - No hazard detection hardware
 - Binary code compatibility

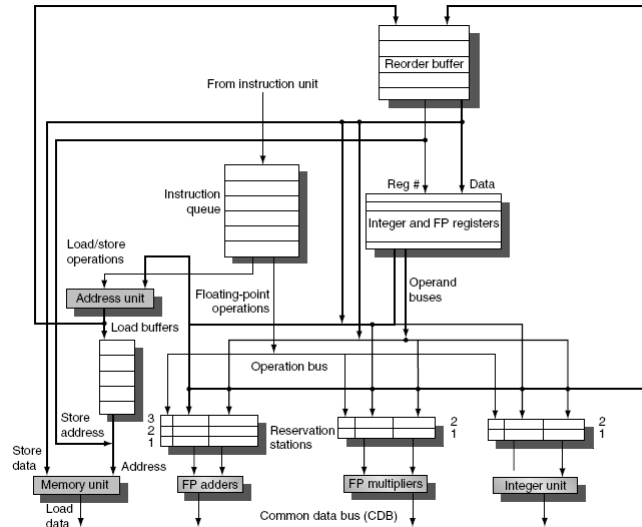
67

Dynamic Scheduling, Multiple Issue, and Speculation

- ▶ Modern microarchitectures:
 - Dynamic scheduling + multiple issue + speculation
- ▶ Two approaches:
 - Assign reservation stations and update pipeline control table in half clock cycles
 - Only supports 2 instructions/clock
 - Design logic to handle any possible dependencies between the instructions
 - Hybrid approaches
- ▶ Issue logic can become bottleneck

68

Overview of Design



69

Multiple Issue

- ▶ Limit the number of instructions of a given class that can be issued in a “bundle”
 - I.e. on FP, one integer, one load, one store
- ▶ Examine all the dependencies among the instructions in the bundle
- ▶ If dependencies exist in bundle, encode them in reservation stations
- ▶ Also need multiple completion/commit

70

Example

```

Loop: LD R2,0(R1)      ;R2=array element
      DADDIU R2,R2,#1  ;increment R2
      SD R2,0(R1)     ;store result
      DADDIU R1,R1,#8  ;increment pointer
      BNE R2,R3,LOOP  ;branch if not last element
  
```

71

Example (No Speculation)

Iteration number	Instructions	Issues at clock cycle number	Executes at clock cycle number	Memory access at clock cycle number	Write CDB at clock cycle number	Comment
1	LD R2,0(R1)	1	2	3	4	First issue
1	DADDIU R2,R2,#1	1	5		6	Wait for LW
1	SD R2,0(R1)	2	3	7		Wait for DADDIU
1	DADDIU R1,R1,#8	2	3		4	Execute directly
1	BNE R2,R3,LOOP	3	7			Wait for DADDIU
2	LD R2,0(R1)	4	8	9	10	Wait for BNE
2	DADDIU R2,R2,#1	4	11		12	Wait for LW
2	SD R2,0(R1)	5	9	13		Wait for DADDIU
2	DADDIU R1,R1,#8	5	8		9	Wait for BNE
2	BNE R2,R3,LOOP	6	13			Wait for DADDIU
3	LD R2,0(R1)	7	14	15	16	Wait for BNE
3	DADDIU R2,R2,#1	7	17		18	Wait for LW
3	SD R2,0(R1)	8	15	19		Wait for DADDIU
3	DADDIU R1,R1,#8	8	14		15	Wait for BNE
3	BNE R2,R3,LOOP	9	19			Wait for DADDIU

72

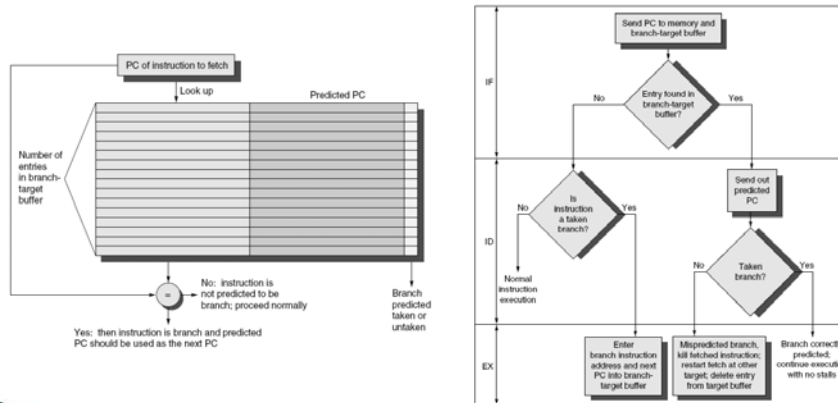
Example

Iteration number	Instructions	Issues at clock number	Executes at clock number	Read access at clock number	Write CDB at clock number	Commits at clock number	Comment
1	LD R2,0(R1)	1	2	3	4	5	First issue
1	DADDIU R2,R2,#1	1	5		6	7	Wait for LW
1	SD R2,0(R1)	2	3			7	Wait for DADDIU
1	DADDIU R1,R1,#8	2	3		4	8	Commit in order
1	BNE R2,R3,LOOP	3	7			8	Wait for DADDIU
2	LD R2,0(R1)	4	5	6	7	9	No execute delay
2	DADDIU R2,R2,#1	4	8			9	Wait for LW
2	SD R2,0(R1)	5	6			10	Wait for DADDIU
2	DADDIU R1,R1,#8	5	6		7	11	Commit in order
2	BNE R2,R3,LOOP	6	10			11	Wait for DADDIU
3	LD R2,0(R1)	7	8	9	10	12	Earliest possible
3	DADDIU R2,R2,#1	7	11		12	13	Wait for LW
3	SD R2,0(R1)	8	9			13	Wait for DADDIU
3	DADDIU R1,R1,#8	8	9		10	14	Executes earlier
3	BNE R2,R3,LOOP	9	13			14	Wait for DADDIU

73

Branch-Target Buffer

- ▶ Need high instruction bandwidth!
 - Branch-Target buffers
 - Next PC prediction buffer, indexed by current PC



74

Branch Folding

- ▶ Optimization:
 - Larger branch-target buffer
 - Add target instruction into buffer to deal with longer decoding time required by larger buffer
 - “Branch folding”

75

Return Address Predictor

- ▶ Most unconditional branches come from function returns
- ▶ The same procedure can be called from multiple sites
 - Causes the buffer to potentially forget about the return address from previous calls
- ▶ Create return address buffer organized as a stack

76

Integrated Instruction Fetch Unit

- ▶ Design monolithic unit that performs:
 - Branch prediction
 - Instruction prefetch
 - Fetch ahead
 - Instruction memory access and buffering
 - Deal with crossing cache lines

77

Register Renaming

- ▶ Register renaming vs. reorder buffers
 - Instead of virtual registers from reservation stations and reorder buffer, create a single register pool
 - Contains visible registers and virtual registers
 - Use hardware-based map to rename registers during issue
 - WAW and WAR hazards are avoided
 - Speculation recovery occurs by copying during commit
 - Still need a ROB-like queue to update table in order
 - Simplifies commit:
 - Record that mapping between architectural register and physical register is no longer speculative
 - Free up physical register used to hold older value
 - In other words: SWAP physical registers on commit
 - Physical register de-allocation is more difficult

78

Integrated Issue and Renaming

- ▶ Combining instruction issue with register renaming:
 - Issue logic pre-reserves enough physical registers for the bundle (fixed number?)
 - Issue logic finds dependencies within bundle, maps registers as necessary
 - Issue logic finds dependencies between current bundle and already in-flight bundles, maps registers as necessary

79

How Much?

- ▶ How much to speculate
 - Mis-speculation degrades performance and power relative to no speculation
 - May cause additional misses (cache, TLB)
 - Prevent speculative code from causing higher costing misses (e.g. L2)
- ▶ Speculating through multiple branches
 - Complicates speculation recovery
 - No processor can resolve multiple branches per cycle

80

Energy Efficiency

- ▶ Speculation and energy efficiency
 - Note: speculation is only energy efficient when it significantly improves performance

- ▶ Value prediction
 - Uses:
 - Loads that load from a constant pool
 - Instruction that produces a value from a small set of values
 - Not been incorporated into modern processors
 - Similar idea--*address aliasing prediction*--is used on some processors