

BLM6112
Advanced Computer Architecture
Instruction-Level Parallelism and Its
Exploitation

Prof. Dr. Nizamettin AYDIN

naydin@yildiz.edu.tr

<http://www3.yildiz.edu.tr/~naydin>

Outline

- Instruction Level Parallelism
- Loop Unrolling
- Instruction Dependencies
- Dynamic Scheduling
- Tomasulo Algorithm
- Speculative Approach
- Dynamic Branch Prediction
- Multiple Instruction Issue

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 PMP processors
 - PMP: Parallel Microprogrammed Processor
 - Compiler-based static approaches
 - Not as successful outside of scientific applications

Instruction-Level Parallelism

- When exploiting Instruction-Level Parallelism (ILP), goal is to maximize CPI (Cycles Per Instruction)
 - 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

Data Dependence

- Loop-Level Parallelism
 - Unroll loop statically or dynamically
 - Use SIMD (vector processors and GPUs)
 - SIMD (Single Instruction Multiple Data)
- 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

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

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 register 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:


```
add x1,x2,x3
beq x4,x0,L
sub x1,x1,x6
L: ...
or x7,x1,x8
```

 - or instruction dependent on add and sub
- Example 2:


```
add x1,x2,x3
beq x12,x0,skip
sub x4,x5,x6
add x5,x4,x9
skip:
or x7,x8,x9
```

 - Assume $x4$ isn't used after skip
 - Possible to move sub before the branch

9

Compiler Techniques for Exposing ILP

1.4.2020

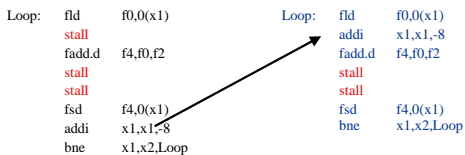
- Pipeline scheduling
 - Separate dependent instruction from the source instruction by the pipeline latency of the source instruction
- Example:


```
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



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

Loop Unrolling

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

```
Loop:  fld    f0,0(x1)
      fadd.d  f4,f0,f2
      fsd    4,0(x1) //drop addi & bne
      fld    f6,-8(x1)
      fadd.d  f8,f6,f2
      fsd    f8,-8(x1) //drop addi & bne
      fld    f0,-16(x1)
      fadd.d  f12,f0,f2
      fsd    f12,-16(x1) //drop addi & bne
      fld    f14,-24(x1)
      fadd.d  f16,f14,f2
      fsd    f16,-24(x1)
      addi   x1,x1,-32
      bne    x1,x2,Loop
```

■ note: number of live registers vs. original loop

12

Loop Unrolling/Pipeline Scheduling

- Pipeline schedule the unrolled loop:

```

Loop: fld    f0,0(x1)
      fld    f6,-8(x1)
      fld    f8,-16(x1)
      fld    f14,-24(x1)
      fadd.d f4,f0,f2
      fadd.d f8,f6,f2
      fadd.d f12,f0,f2
      fadd.d f16,f14,f2
      fsd    f4,0(x1)
      fsd    f8,-8(x1)
      fsd    f12,-16(x1)
      fsd    f16,-24(x1)
      addi   x1,x1,-32
      bne   x1,x2,Loop
    
```

- 14 cycles
- 3.5 cycles per element

13

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”

14

FP Loop: Where are the Hazards?

- First translate into MIPS code:
 - To simplify, assume 8 is lowest address, R1 = base address of X and F2 = s

```

Loop: L.D    F0,0(R1);F0=vector element
      ADD.D  F4,F0,F2;add scalar in F2
      S.D    0(R1),F4;store result
      DADDUI R1,R1,-8;decrement pointer 8B (DW)
      BNEZ  R1,Loop ;branch R1!=zero
    
```

15

FP Loop Showing Stalls

```

1 Loop: L.D    F0,0(R1) ;F0=vector element
2      stall
3      ADD.D  F4,F0,F2 ;add scalar in F2
4      stall
5      stall
6      S.D    0(R1),F4 ;store result
7      DADDUI R1,R1,-8 ;decrement pointer 8B (DW)
8      stall ;assumes can't forward to branch
9      BNEZ  R1,Loop ;branch R1!=zero
    
```

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

- 9 clock cycles: Rewrite code to minimize stalls?

16

Revised FP Loop Minimizing Stalls

```

1 Loop: L.D    F0,0(R1)
2      DADDUI R1,R1,-8
3      ADD.D  F4,F0,F2
4      stall
5      stall
6      S.D    8(R1),F4 ;altered offset when move DADDUI
7      BNEZ  R1,Loop
    
```

Swap DADDUI and S.D by changing address of S.D

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

7 clock cycles, but just 3 for execution (L.D, ADD.D,S.D), 4 for loop overhead; How to make faster?

17

Unroll Loop Four Times (straightforward way)

```

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

Rewrite loop to minimize stalls?

27 clock cycles, or 6.75 per iteration (Assumes R1 is multiple of 4)

18

Unrolled Loop That Minimizes Stalls

```

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     S.D    8(R1), F16    ; 8-32 = -24
14     BNEZ  R1, LOOP
    
```

14 clock cycles, or 3.5 per iteration

19

Unrolled Loop Detail

- Do not usually know upper bound of loop
- Suppose it is n , and we would like to unroll the loop to make k copies of the body
- Instead of a single unrolled loop, we generate a pair of consecutive loops:
 - 1st executes $(n \bmod k)$ times and has a body that is the original loop
 - 2nd is the unrolled body surrounded by an outer loop that iterates (n/k) times
- For large values of n , most of the execution time will be spent in the unrolled loop

20

5 Loop Unrolling Decisions

- Requires understanding how one instruction depends on another and how the instructions can be changed or reordered given the dependences:
 - Determine loop unrolling useful by finding that loop iterations were independent (except for maintenance code)
 - Use different registers to avoid unnecessary constraints forced by using same registers for different computations
 - Eliminate the extra test and branch instructions and adjust the loop termination and iteration code
 - Determine that loads and stores in unrolled loop can be interchanged by observing that loads and stores from different iterations are independent
 - Transformation requires analyzing memory addresses and finding that they do not refer to the same address
 - Schedule the code, preserving any dependences needed to yield the same result as the original code

21

In-class Exercise

- Identify data hazards in the code below:
 - MULTD F3, F4, F2
 - ADDD F2, F6, F1
 - SD F2, 0(F3)
- For each of the following code fragments, identify each type of dependence that a compiler will find (a fragment may have no dependences) and whether a compiler could schedule the two instructions (i.e., change their orders).

1. DADDI R1, R1, #4	2. DADD R3, R1, R2
LD R2, 7(R1)	SD R2, 7(R1)
3. SD R2, 7(R1)	4. BEZ R1, place
SD F2, 200(R7)	SD R1, 7(R1)

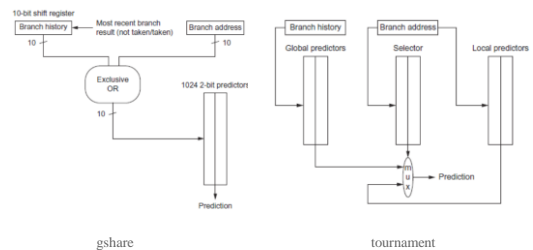
22

Branch Prediction

- Basic 2-bit predictor:
 - For each branch:
 - Predict taken or not taken
 - If the prediction is wrong two consecutive times, change prediction
- Correlating predictor:
 - Multiple 2-bit predictors for each branch
 - One for each possible combination of outcomes of preceding n branches
 - (m, n) predictor: behavior from last m branches to choose from 2^m n -bit predictors
- Tournament predictor:
 - Combine correlating predictor with local predictor

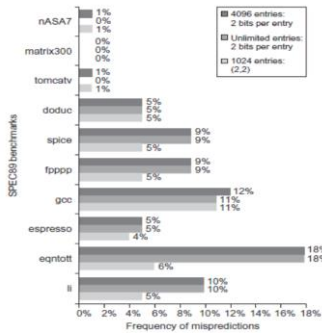
23

Branch Prediction



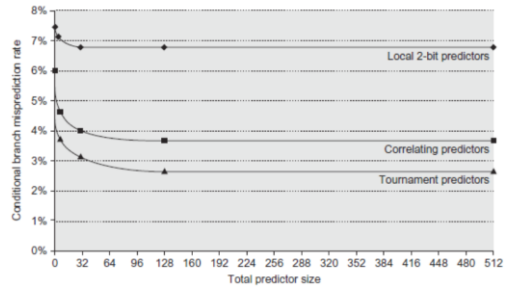
24

Branch Prediction Performance



25

Branch Prediction Performance

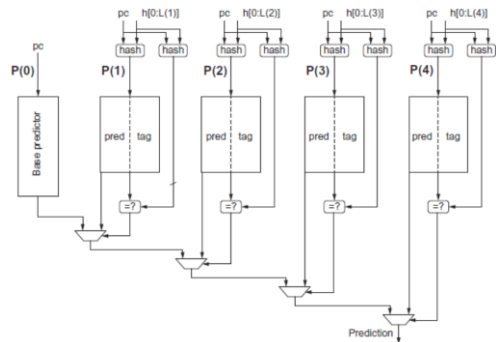


26

Tagged Hybrid Predictors

- Need to have predictor for each branch and history
 - Problem: this implies huge tables
 - Solution:
 - Use hash tables, whose hash value is based on branch address and branch history
 - Longer histories may lead to increased chance of hash collision, so use multiple tables with increasingly shorter histories

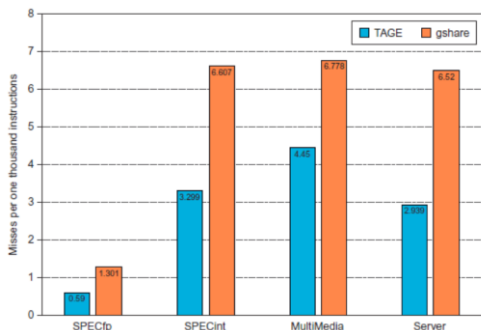
Tagged Hybrid Predictors



27

28

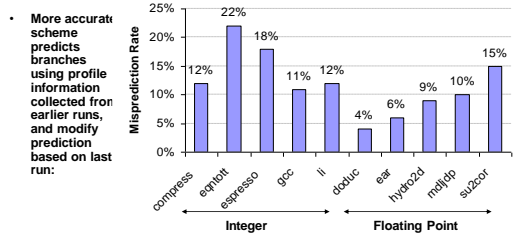
Tagged Hybrid Predictors



29

Static Branch Prediction

- Previous lectures showed scheduling code around delayed branch
- To reorder code around branches, need to predict branch statically when compile
- Simplest scheme is to predict a branch as taken
 - Average misprediction = untaken branch frequency = 34% SPEC



30

Dynamic Branch Prediction

- Why does prediction work?
 - Underlying algorithm has regularities
 - Data that is being operated on has regularities
 - Instruction sequence has redundancies that are artifacts of way that humans/compiler think about problems
- Is dynamic branch prediction better than static branch prediction?
 - Seems to be
 - There are a small number of important branches in programs which have dynamic behavior

31

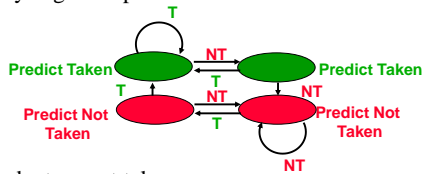
Dynamic Branch Prediction

- Performance = $f(\text{accuracy}, \text{cost of misprediction})$
- Branch History Table: Lower bits of PC address index table of 1-bit values
 - Says whether or not branch taken last time
 - No address check
- Problem: in a loop, 1-bit BHT will cause two mispredictions (avg is 9 iterations before exit):
 - End of loop case, when it exits instead of looping as before
 - First time through loop on next time through code, when it predicts exit instead of looping

32

Dynamic Branch Prediction

- Solution: 2-bit scheme where change prediction only if get misprediction twice

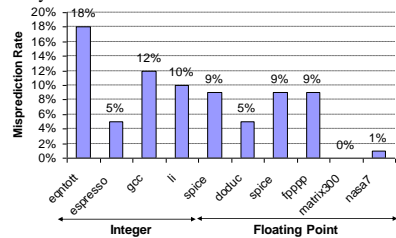


- Red: stop, not taken
- Green: go, taken
- Adds *hysteresis* to decision making process

33

BHT Accuracy

- Mispredict because either:
 - Wrong guess for that branch
 - Got branch history of wrong branch when index the table
- 4096 entry table:



34

Correlated Branch Prediction

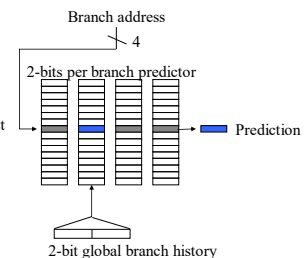
- Idea: record m most recently executed branches as taken or not taken, and use that pattern to select the proper n -bit branch history table
- In general, a (m,n) predictor means recording last m branches to select between 2^m history tables, each with n -bit counters
 - Thus, old 2-bit BHT is a $(0,2)$ predictor
- Global Branch History: m -bit shift register keeping T/NT status of last m branches.
- Each entry in table has 2^m n -bit predictors.

35

Correlating Branches

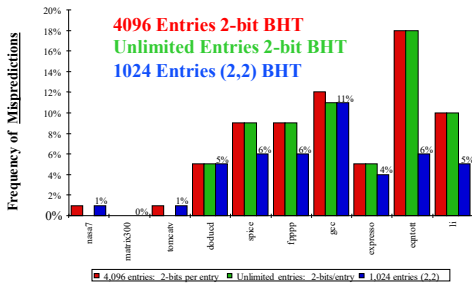
(2,2) predictor

- Behavior of recent branches selects between four predictions of next branch, updating just that prediction



36

Accuracy of Different Schemes



37

Tournament Predictors

Tournament predictor using, say, 4K 2-bit counters indexed by local branch address.

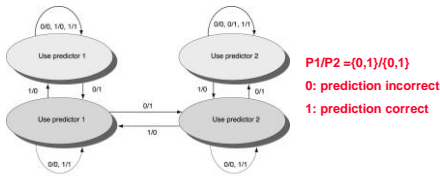
Chooses between:

- Global predictor
 - 4K entries index by history of last 12 branches ($2^{12} = 4K$)
 - Each entry is a standard 2-bit predictor
- Local predictor
 - Local history table: 1024 10-bit entries recording last 10 branches, index by branch address
 - The pattern of the last 10 occurrences of that particular branch used to index table of 1K entries with 3-bit saturating counters

38

Tournament Predictors

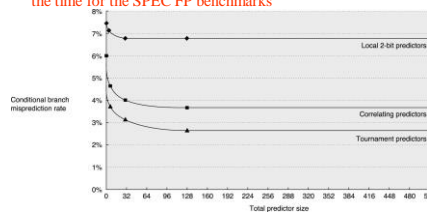
- Multilevel branch predictor
- Use n -bit saturating counter to choose between [predictors](#)
- Usual choice between global and local predictors



39

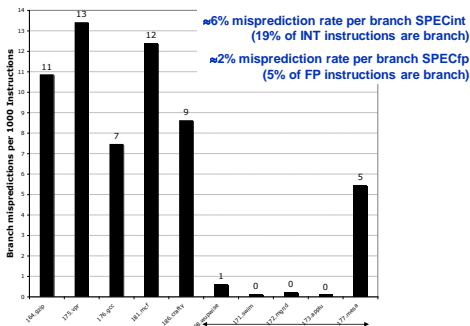
Comparing Predictors

- Advantage of tournament predictor is ability to select the right predictor for a particular branch
 - Particularly crucial for integer benchmarks.
 - A typical tournament predictor will select the global predictor almost 40% of the time for the SPEC integer benchmarks and less than 15% of the time for the SPEC FP benchmarks



40

Pentium 4 Misprediction Rate (per 1000 instructions, not per branch)



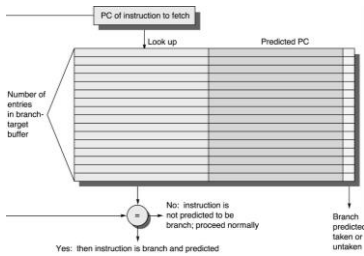
41

Branch Target Buffers (BTB)

- Branch target calculation is costly and stalls the instruction fetch.
- BTB stores PCs the same way as caches
- The PC of a branch is sent to the BTB
- When a match is found the corresponding Predicted PC is returned
- If the branch was predicted taken, instruction fetch continues at the returned predicted PC

42

Branch Target Buffers



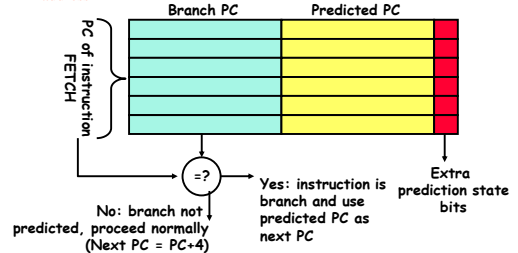
Branch target folding: for unconditional branches store the target instructions themselves in the buffer!

43

Need Address at Same Time as Prediction

- Branch Target Buffer (BTB): Address of branch index to get prediction AND branch address (if taken)

– Note: must check for branch match now, since can't use wrong branch address



44

Dynamic Branch Prediction Summary

- Prediction becoming important part of execution
- Branch History Table: 2 bits for loop accuracy
- Correlation: Recently executed branches correlated with next branch
 - Either different branches (GA)
 - Or different executions of same branches (PA)
- Tournament predictors take insight to next level, by using multiple predictors
 - usually one based on global information and one based on local information, and combining them with a selector
 - In 2006, tournament predictors using $\approx 30K$ bits are in processors like the Power5 and Pentium 4
- Branch Target Buffer: include branch address & prediction;
- Branch target folding.

45

Dynamic Scheduling

- Rearrange order of instructions to reduce stalls while maintaining data flow
- Advantages:
 - Compiler doesn't need to have knowledge of microarchitecture
 - Handles cases where dependencies are unknown at compile time
- Disadvantage:
 - Substantial increase in hardware complexity
 - Complicates exceptions

46

Dynamic Scheduling

- Dynamic scheduling implies:
 - Out-of-order execution
 - Out-of-order completion
- Example 1:


```
fdiv.d  f0,f2,f4
fadd.d  f10,f0,f8
fsub.d  f12,f8,f14
```

 - fsub.d is not dependent, issue before fadd.d

47

Dynamic Scheduling

- Example 2:


```
fdiv.d  f0,f2,f4
fmul.d  f6,f0,f8
fadd.d  f0,f10,f14
```

 - fadd.d is not dependent, but the antidependence makes it impossible to issue earlier without register renaming

48

Register Renaming

- Example 3:

```

fdiv.d    f0,f2,f4
fadd.d    f6,f0,f8
fsd       f6,0(x1)
fsub.d    f8,f10,f14
fmul.d    f6,f10,f8
    
```

Antidependence f8
Output dependence f6

– name dependence with f0,f6,f8

49

Register Renaming

- Example 3:

```

fdiv.d    f0,f2,f4
fadd.d    S,f0,f8
fsd       S,0(x1)
fsub.d    T,f10,f14
fmul.d    f6,f10,T
    
```

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

50

Register Renaming

- Tomasulo's Approach
 - Tracks when operands are available
 - Introduces register renaming in hardware
 - Minimizes WAW and WAR hazards
- 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

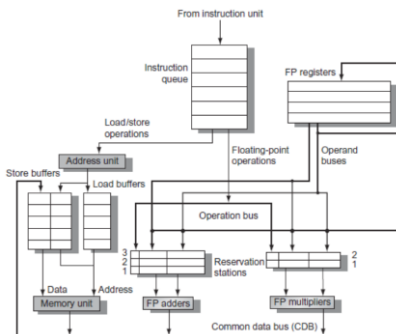
51

Register Renaming

- 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
- Load and store buffers
 - Contain data and addresses, act like reservation stations

52

Tomasulo's Algorithm



53

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)

54

Example

Instruction	Instruction status		
	Issue	Execute	Write result
Fld f6, 32(x2)	√	√	√
Fld f2, 44(x3)	√	√	
Fmul.d f0, f2, f4	√		
Fsub.d f0, f2, f6	√		
Fdiv.d f0, f0, f6	√		
Fadd.d f6, f8, f2	√		

Reservation stations							
Name	Busy	Op	Vj	Vk	Qj	Qk	A
Load1	No						
Load2	Yes	Load					44 + Regs[x3]
Add1	Yes	SUB		Mem[32 + Regs[x2]]		Load2	
Add2	Yes	ADD				Add1	Load2
Add3	No						
Mult1	Yes	MUL		Regs[f4]		Load2	
Mult2	Yes	DIV		Mem[32 + Regs[x2]]		Mult1	

Register status										
Field	f0	f2	f4	f6	f8	f10	f12	...	f30	
Qj	Mult1	Load2		Add2	Add1	Mult2				

55

Tomasulo's Algorithm

- Example loop:

```

Loop: fld    f0,0(x1)
      fmul.d f4,f0,f2
      fsd    f4,0(x1)
      addi   x1,x1,8
      bne   x1,x2,Loop // branches if x1 != x2
    
```

56

Tomasulo's Algorithm

Instruction	Instruction status			
	From iteration	Issue	Execute	Write result
Fld f0, 0(x1)	1	√	√	
Fmul.d f4, f0, f2	1	√		
Fsd f4, 0(x1)	1	√		
Fld f0, 0(x1)	2	√	√	
Fmul.d f4, f0, f2	2	√		
Fsd f4, 0(x1)	2	√		

Reservation stations							
Name	Busy	Op	Vj	Vk	Qj	Qk	A
Load1	Yes	Load					Regs[x1] + 0
Load2	Yes	Load					Regs[x1] - 8
Add1	No						
Add2	No						
Add3	No						
Mult1	Yes	MUL		Regs[f2]		Load1	
Mult2	Yes	MUL		Regs[f2]		Load2	
Store1	Yes	Store		Regs[x1]			Mult1
Store2	Yes	Store		Regs[x1] - 8			Mult2

Register status										
Field	f0	f2	f4	f6	f8	f10	f12	...	f30	
Qj		Load2		Mult2						

57

Dynamic Scheduling

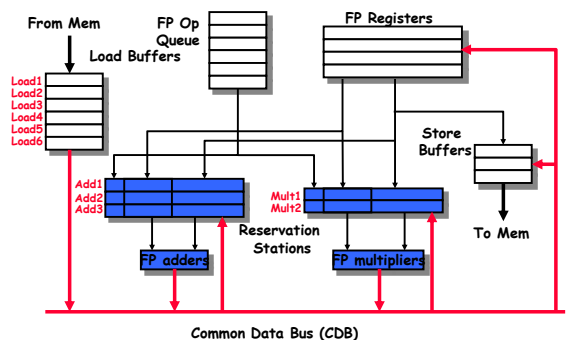
- Simple pipeline had 1 stage to check both structural and data hazards: Instruction Decode (ID), also called Instruction Issue
- Split the ID pipe stage of simple 5-stage pipeline into 2 stages:
- **Issue**—Decode instructions, check for structural hazards
- **Read operands**—Wait until no data hazards, then read operands

58

Tomasulo Algorithm

- Control & buffers distributed with Function Units (FU)
 - FU buffers called **Reservation Stations (RS)**; have pending operands
- Registers in instructions replaced by values or pointers to RS;
 - form of **register renaming**;
 - avoids WAR and WAW hazards
 - More reservation stations than registers, so can do optimizations compilers can't
- Results to FU from RS, **not through registers**, over **Common Data Bus** that broadcasts results to all FUs
- Load and Stores treated as FUs with RSs as well
- Integer instructions can go past branches, allowing FP ops beyond basic block in FP queue

Tomasulo Organization



59

60

Reservation Station Components

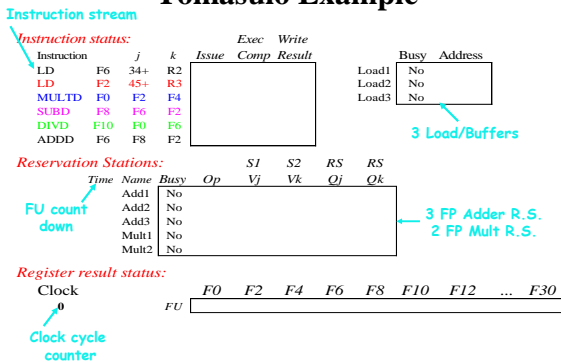
- Op: Operation to perform in the unit (e.g., + or -)
- Vj, Vk: Value of Source operands
 - Store buffers has V field, result to be stored
- Qj, Qk: Reservation stations producing source registers (value to be written)
 - Note: Qj,Qk=0 => ready
 - Store buffers only have Qj for RS producing result
- Busy: Indicates reservation station or FU is busy

Register result status—Indicates which functional unit will write each register, if one exists. Blank when no pending instructions that will write that register.

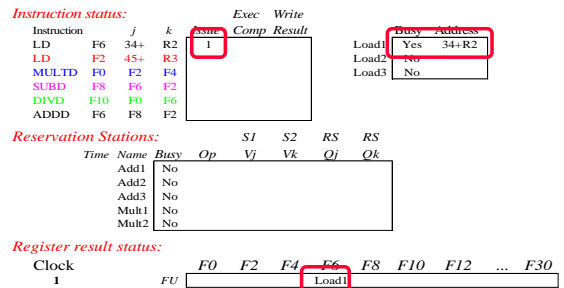
Three Stages of Tomasulo Algorithm

- Issue**—get instruction from FP Op Queue
 - If reservation station free (no structural hazard), control issues instr & sends operands (renames registers).
- Execute**—operate on operands (EX)
 - When both operands ready then execute; if not ready, watch Common Data Bus for result
- Write result**—finish execution (WB)
 - Write on Common Data Bus to all awaiting units; mark reservation station available
 - Normal data bus: data + destination (“go to” bus)
 - Common data bus: data + source (“come from” bus)
 - 64 bits of data + 4 bits of Functional Unit source address
 - Write if matches expected Functional Unit (produces result)
 - Does the broadcast
 - Example speed:
 - 3 clocks for F1 .pt. +,-; 10 for *; 40 clks for /

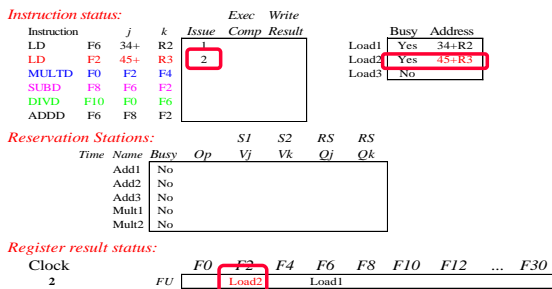
Tomasulo Example



Tomasulo Example Cycle 1

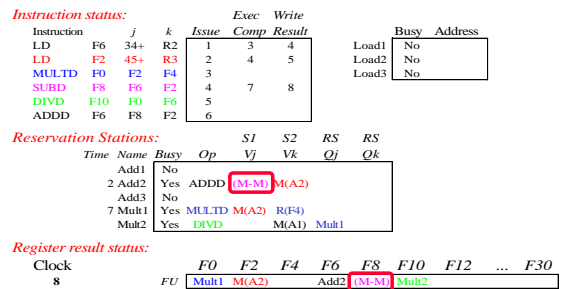


Tomasulo Example Cycle 2



Note: Can have multiple loads outstanding

Tomasulo Example Cycle 8



Tomasulo Example Cycle 3

Instruction status:

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

Reservation Stations:

Time	Name	Busy	Op	Vj	Vk	Qj	Qk
0	Add1	No					
0	Add2	No					
0	Add3	No					
1	Mult1	Yes	MULTD		R(F4)	Load2	
2	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?

67

Tomasulo Example Cycle 4

Instruction status:

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

Reservation Stations:

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

Register result status:

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

- Load2 completing; what is waiting for Load2?

68

Tomasulo Example Cycle 5

Instruction status:

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

Reservation Stations:

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

Register result status:

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

- Timer starts down for Add1, Mult1

69

Tomasulo Example Cycle 6

Instruction status:

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

Reservation Stations:

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

Register result status:

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

- Issue ADDD here despite name dependency on F6?

70

Tomasulo Example Cycle 7

Instruction status:

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

Reservation Stations:

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

Register result status:

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

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

71

Tomasulo Example Cycle 9

Instruction status:

Instruction	j	k	Issue	Comp	Result	Busy	Address
LD	F6	34+	R2	1	3	Load1	No
LD	F2	45+	R3	2	4	Load2	No
MULTD	F0	F2	F4	3		Load3	No
SUBD	F8	F6	F2	4	7		
DIVD	F10	F0	F6	5			
ADDD	F6	F8	F2	6			

Reservation Stations:

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

Register result status:

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

72

Tomasulo Example Cycle 10

Instruction status:

Instruction	j	k	Exec			Write	Busy	Address
			Issue	Comp	Result			
LD	F6	34+	R2	1	3	4	No	No
LD	F2	45+	R3	2	4	5	No	No
MULTD	F0	F2	F4	3			No	No
SUBD	F8	F6	F2	4	7	8		
DIVD	F10	F0	F6	5				
ADDD	F6	F8	F2	6	10	11		

Reservation Stations:

Time	Name	Busy	Op	Vj	Vk	Rs	Rd
0	Add1	No					
0	Add2	Yes	ADDD	(M-M)	M(A2)		
0	Add3	No					
5	Mult1	Yes	MULTD	M(A2)	R(F4)		
5	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		

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

73

Tomasulo Example Cycle 11

Instruction status:

Instruction	j	k	Exec			Write	Busy	Address
			Issue	Comp	Result			
LD	F6	34+	R2	1	3	4	No	No
LD	F2	45+	R3	2	4	5	No	No
MULTD	F0	F2	F4	3			No	No
SUBD	F8	F6	F2	4	7	8		
DIVD	F10	F0	F6	5				
ADDD	F6	F8	F2	6	10	11		

Reservation Stations:

Time	Name	Busy	Op	Vj	Vk	Rs	Rd
0	Add1	No					
0	Add2	No					
0	Add3	No					
4	Mult1	Yes	MULTD	M(A2)	R(F4)		
5	Mult2	Yes	DIVD	M(A1)	Mult1		

Register result status:

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

- Write result of ADDD here?
- All quick instructions complete in this cycle!

74

Tomasulo Example Cycle 12

Instruction status:

Instruction	j	k	Exec			Write	Busy	Address
			Issue	Comp	Result			
LD	F6	34+	R2	1	3	4	No	No
LD	F2	45+	R3	2	4	5	No	No
MULTD	F0	F2	F4	3			No	No
SUBD	F8	F6	F2	4	7	8		
DIVD	F10	F0	F6	5				
ADDD	F6	F8	F2	6	10	11		

Reservation Stations:

Time	Name	Busy	Op	Vj	Vk	Rs	Rd
0	Add1	No					
0	Add2	No					
0	Add3	No					
3	Mult1	Yes	MULTD	M(A2)	R(F4)		
5	Mult2	Yes	DIVD	M(A1)	Mult1		

Register result status:

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

75

Tomasulo Example Cycle 13

Instruction status:

Instruction	j	k	Exec			Write	Busy	Address
			Issue	Comp	Result			
LD	F6	34+	R2	1	3	4	No	No
LD	F2	45+	R3	2	4	5	No	No
MULTD	F0	F2	F4	3			No	No
SUBD	F8	F6	F2	4	7	8		
DIVD	F10	F0	F6	5				
ADDD	F6	F8	F2	6	10	11		

Reservation Stations:

Time	Name	Busy	Op	Vj	Vk	Rs	Rd
0	Add1	No					
0	Add2	No					
0	Add3	No					
2	Mult1	Yes	MULTD	M(A2)	R(F4)		
5	Mult2	Yes	DIVD	M(A1)	Mult1		

Register result status:

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

76

Tomasulo Example Cycle 14

Instruction status:

Instruction	j	k	Exec			Write	Busy	Address
			Issue	Comp	Result			
LD	F6	34+	R2	1	3	4	No	No
LD	F2	45+	R3	2	4	5	No	No
MULTD	F0	F2	F4	3	15		No	No
SUBD	F8	F6	F2	4	7	8		
DIVD	F10	F0	F6	5				
ADDD	F6	F8	F2	6	10	11		

Reservation Stations:

Time	Name	Busy	Op	Vj	Vk	Rs	Rd
0	Add1	No					
0	Add2	No					
0	Add3	No					
1	Mult1	Yes	MULTD	M(A2)	R(F4)		
5	Mult2	Yes	DIVD	M(A1)	Mult1		

Register result status:

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

77

Tomasulo Example Cycle 15

Instruction status:

Instruction	j	k	Exec			Write	Busy	Address
			Issue	Comp	Result			
LD	F6	34+	R2	1	3	4	No	No
LD	F2	45+	R3	2	4	5	No	No
MULTD	F0	F2	F4	3	15		No	No
SUBD	F8	F6	F2	4	7	8		
DIVD	F10	F0	F6	5				
ADDD	F6	F8	F2	6	10	11		

Reservation Stations:

Time	Name	Busy	Op	Vj	Vk	Rs	Rd
0	Add1	No					
0	Add2	No					
0	Add3	No					
0	Mult1	Yes	MULTD	M(A2)	R(F4)		
5	Mult2	Yes	DIVD	M(A1)	Mult1		

Register result status:

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

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

78

Tomasulo Example Cycle 16

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	15	16		No		
SUBD	F8	F6	F2	4	7	8		No		
DIVD	F10	F0	F6	5				No		
ADDD	F6	F8	F2	6	10	11		No		

Reservation Stations:

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

Register result status:

Clock	F0	F2	F4	F6	F8	F10	F12	...	F30
16									
	FU	M*F4	M(A2)	(M-M+N)	(M-M)	Mult2			

**Faster than light computation
(skip a couple of cycles)**

- Just waiting for Mult2 (DIVD) to complete

79

80

Tomasulo Example Cycle 55

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	15	16		No		
SUBD	F8	F6	F2	4	7	8		No		
DIVD	F10	F0	F6	5				No		
ADDD	F6	F8	F2	6	10	11		No		

Reservation Stations:

Time	Name	Busy	Op	Vj	Vk	Qj	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
55									
	FU	M*F4	M(A2)	(M-M+N)	(M-M)	Mult2			

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	15	16		No		
SUBD	F8	F6	F2	4	7	8		No		
DIVD	F10	F0	F6	5	56			No		
ADDD	F6	F8	F2	6	10	11		No		

Reservation Stations:

Time	Name	Busy	Op	Vj	Vk	Qj	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
56									
	FU	M*F4	M(A2)	(M-M+N)	(M-M)	Mult2			

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

81

82

Tomasulo Example Cycle 57

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	15	16		No		
SUBD	F8	F6	F2	4	7	8		No		
DIVD	F10	F0	F6	5	56	57		No		
ADDD	F6	F8	F2	6	10	11		No		

Reservation Stations:

Time	Name	Busy	Op	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
56									
	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 Loop Example

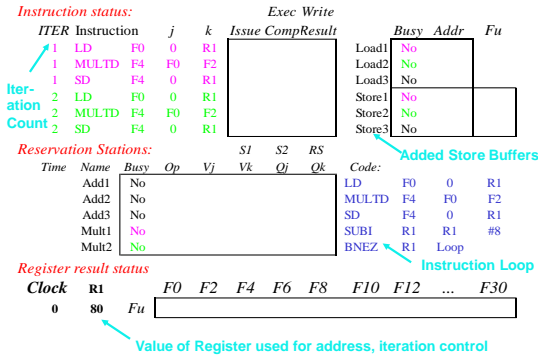
Loop:	LD	F0	0	R1
	MULTD	F4	F0	F2
	SD	F4	0	R1
	SUBI	R1	R1	#8
	BNEZ	R1	Loop	

- This time assume Multiply takes 4 clocks
- Assume 1st load takes 8 clocks (L1 cache miss), 2nd load takes 1 clock (hit)
- To be clear, will show clocks for SUBI, BNEZ
 - Reality: integer instructions ahead of Fl. Pt. Instructions
- Show 2 iterations

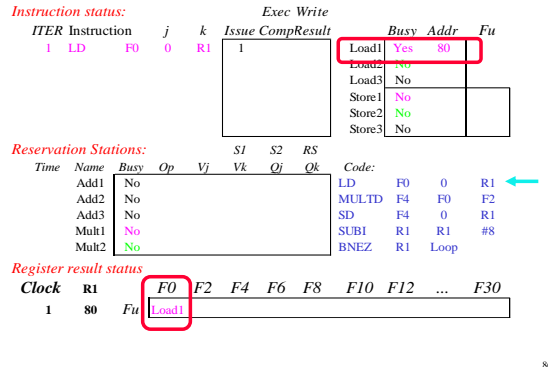
83

84

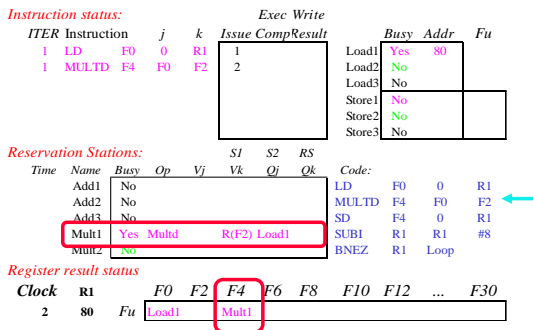
Loop Example



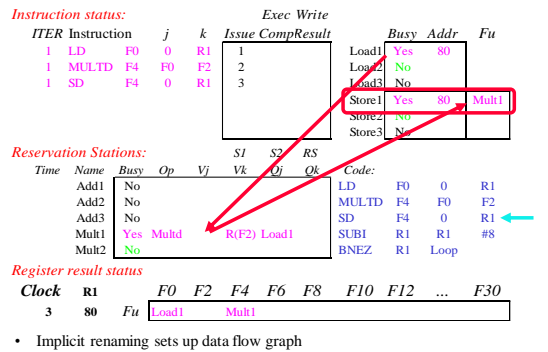
Loop Example Cycle 1



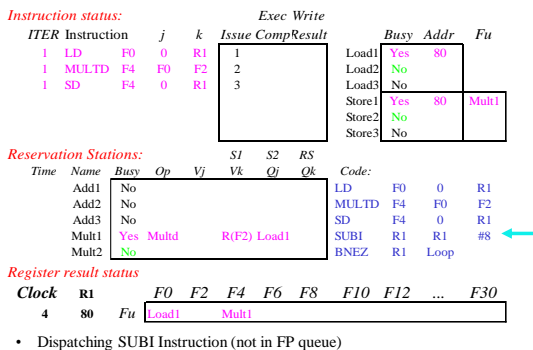
Loop Example Cycle 2



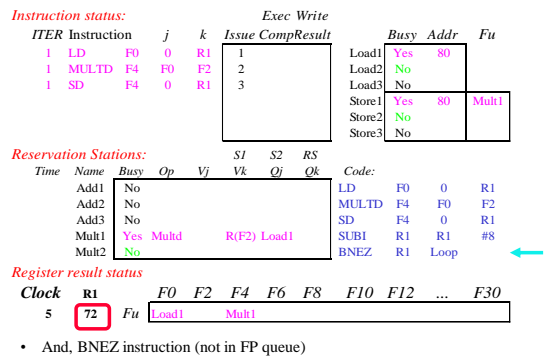
Loop Example Cycle 3



Loop Example Cycle 4



Loop Example Cycle 5



Loop Example Cycle 12

Instruction status:

ITER	Instruction	j	k	Issue	Comp	Result	Busy	Addr	Fu
1	LD	F0	0	R1	1	9	10	Load1	No
1	MULTD	F4	F0	F2	2			Load2	No
1	SD	F4	0	R1	3		64	Load3	Yes
2	LD	F0	0	R1	6	10	11	Store1	Yes 80
2	MULTD	F4	F0	F2	7			Store2	Yes 72
2	SD	F4	0	R1	8			Store3	No

Reservation Stations:

Time	Name	Busy	Op	Vj	Vk	Qj	Qk	Code:
Add1	No							LD F0 0 R1
Add2	No							MULTD F4 F0 F2
Add3	No							SD F4 0 R1
2	Mult1	Yes	Mulhd	M[80]	R(F2)			SUB1 R1 R1 #8
3	Mult2	Yes	Mulhd	M[72]	R(F2)			BNEZ R1 Loop

Register result status

Clock	R1	F0	F2	F4	F6	F8	F10	F12	...	F30
12	64	Fu	Load3	Mult2						

- Why not issue third multiply?

97

Loop Example Cycle 13

Instruction status:

ITER	Instruction	j	k	Issue	Comp	Result	Busy	Addr	Fu
1	LD	F0	0	R1	1	9	10	Load1	No
1	MULTD	F4	F0	F2	2			Load2	No
1	SD	F4	0	R1	3		64	Load3	Yes
2	LD	F0	0	R1	6	10	11	Store1	Yes 80
2	MULTD	F4	F0	F2	7			Store2	Yes 72
2	SD	F4	0	R1	8			Store3	No

Reservation Stations:

Time	Name	Busy	Op	Vj	Vk	Qj	Qk	Code:
Add1	No							LD F0 0 R1
Add2	No							MULTD F4 F0 F2
Add3	No							SD F4 0 R1
1	Mult1	Yes	Mulhd	M[80]	R(F2)			SUB1 R1 R1 #8
2	Mult2	Yes	Mulhd	M[72]	R(F2)			BNEZ R1 Loop

Register result status

Clock	R1	F0	F2	F4	F6	F8	F10	F12	...	F30
13	64	Fu	Load3	Mult2						

- Why not issue third store?

98

Loop Example Cycle 14

Instruction status:

ITER	Instruction	j	k	Issue	Comp	Result	Busy	Addr	Fu
1	LD	F0	0	R1	1	9	10	Load1	No
1	MULTD	F4	F0	F2	2	14		Load2	No
1	SD	F4	0	R1	3		64	Load3	Yes
2	LD	F0	0	R1	6	10	11	Store1	Yes 80
2	MULTD	F4	F0	F2	7			Store2	Yes 72
2	SD	F4	0	R1	8			Store3	No

Reservation Stations:

Time	Name	Busy	Op	Vj	Vk	Qj	Qk	Code:
Add1	No							LD F0 0 R1
Add2	No							MULTD F4 F0 F2
Add3	No							SD F4 0 R1
0	Mult1	Yes	Mulhd	M[80]	R(F2)			SUB1 R1 R1 #8
1	Mult2	Yes	Mulhd	M[72]	R(F2)			BNEZ R1 Loop

Register result status

Clock	R1	F0	F2	F4	F6	F8	F10	F12	...	F30
14	64	Fu	Load3	Mult2						

- Mult1 completing. Who is waiting?

99

Loop Example Cycle 15

Instruction status:

ITER	Instruction	j	k	Issue	Comp	Result	Busy	Addr	Fu
1	LD	F0	0	R1	1	9	10	Load1	No
1	MULTD	F4	F0	F2	2	14	15	Load2	No
1	SD	F4	0	R1	3		64	Load3	Yes
2	LD	F0	0	R1	6	10	11	Store1	Yes 80
2	MULTD	F4	F0	F2	7			Store2	Yes 72
2	SD	F4	0	R1	8			Store3	No

Reservation Stations:

Time	Name	Busy	Op	Vj	Vk	Qj	Qk	Code:
Add1	No							LD F0 0 R1
Add2	No							MULTD F4 F0 F2
Add3	No							SD F4 0 R1
0	Mult1	No						SUB1 R1 R1 #8
0	Mult2	Yes	Mulhd	M[72]	R(F2)			BNEZ R1 Loop

Register result status

Clock	R1	F0	F2	F4	F6	F8	F10	F12	...	F30
15	64	Fu	Load3	Mult2						

- Mult2 completing. Who is waiting?

100

Loop Example Cycle 16

Instruction status:

ITER	Instruction	j	k	Issue	Comp	Result	Busy	Addr	Fu
1	LD	F0	0	R1	1	9	10	Load1	No
1	MULTD	F4	F0	F2	2	14	15	Load2	No
1	SD	F4	0	R1	3		64	Load3	Yes
2	LD	F0	0	R1	6	10	11	Store1	Yes 80
2	MULTD	F4	F0	F2	7	15	16	Store2	Yes 72
2	SD	F4	0	R1	8			Store3	No

Reservation Stations:

Time	Name	Busy	Op	Vj	Vk	Qj	Qk	Code:
Add1	No							LD F0 0 R1
Add2	No							MULTD F4 F0 F2
Add3	No							SD F4 0 R1
4	Mult1	Yes	Mulhd	R(F2)	Load3			SUB1 R1 R1 #8
4	Mult2	No						BNEZ R1 Loop

Register result status

Clock	R1	F0	F2	F4	F6	F8	F10	F12	...	F30
16	64	Fu	Load3	Mult1						

- Register result status

101

Loop Example Cycle 17

Instruction status:

ITER	Instruction	j	k	Issue	Comp	Result	Busy	Addr	Fu
1	LD	F0	0	R1	1	9	10	Load1	No
1	MULTD	F4	F0	F2	2	14	15	Load2	No
1	SD	F4	0	R1	3		64	Load3	Yes
2	LD	F0	0	R1	6	10	11	Store1	Yes 80
2	MULTD	F4	F0	F2	7	15	16	Store2	Yes 72
2	SD	F4	0	R1	8			Store3	Yes 64

Reservation Stations:

Time	Name	Busy	Op	Vj	Vk	Qj	Qk	Code:
Add1	No							LD F0 0 R1
Add2	No							MULTD F4 F0 F2
Add3	No							SD F4 0 R1
Mult1	Yes	Mulhd	R(F2)	Load3				SUB1 R1 R1 #8
Mult2	No							BNEZ R1 Loop

Register result status

Clock	R1	F0	F2	F4	F6	F8	F10	F12	...	F30
17	64	Fu	Load3	Mult1						

- Register result status

102

Loop Example Cycle 18

Instruction status:

ITER	Instruction	j	k	Issue	Comp	Result	Busy	Addr	Fu
1	LD	F0	0	R1	1	9	10	Load1	No
1	MULTD	F4	F0	F2	2	14	15	Load2	No
1	SD	F4	0	R1	3	18	19	Load3	Yes 64
2	LD	F0	0	R1	6	10	11	Store1	Yes 80 [80]*R3
2	MULTD	F4	F0	F2	7	15	16	Store2	Yes 72 [72]*R3
2	SD	F4	0	R1	8	19	20	Store3	Yes 64 Multi1

Reservation Stations:

Time	Name	Busy	Op	Vj	Vk	Qj	Qk	Code:
Add1	No							LD F0 0 R1
Add2	No							MULTD F4 F0 F2
Add3	No							SD F4 0 R1
Mult1	Yes	Multd				R(F2)	Load3	SUBI R1 R1 #8
Mult2	No							BNEZ R1 Loop

Register result status

Clock	R1	F0	F2	F4	F6	F8	F10	F12	...	F30
18	64	Fu	Load3	Multi1						

103

Loop Example Cycle 19

Instruction status:

ITER	Instruction	j	k	Issue	Comp	Result	Busy	Addr	Fu
1	LD	F0	0	R1	1	9	10	Load1	No
1	MULTD	F4	F0	F2	2	14	15	Load2	No
1	SD	F4	0	R1	3	18	19	Load3	Yes 64
2	LD	F0	0	R1	6	10	11	Store1	No
2	MULTD	F4	F0	F2	7	15	16	Store2	Yes 72 [72]*R3
2	SD	F4	0	R1	8	19	20	Store3	Yes 64 Multi1

Reservation Stations:

Time	Name	Busy	Op	Vj	Vk	Qj	Qk	Code:
Add1	No							LD F0 0 R1
Add2	No							MULTD F4 F0 F2
Add3	No							SD F4 0 R1
Mult1	Yes	Multd				R(F2)	Load3	SUBI R1 R1 #8
Mult2	No							BNEZ R1 Loop

Register result status

Clock	R1	F0	F2	F4	F6	F8	F10	F12	...	F30
19	56	Fu	Load3	Multi1						

104

Loop Example Cycle 20

Instruction status:

ITER	Instruction	j	k	Issue	Comp	Result	Busy	Addr	Fu
1	LD	F0	0	R1	1	9	10	Load1	Yes 36
1	MULTD	F4	F0	F2	2	14	15	Load2	No
1	SD	F4	0	R1	3	18	19	Load3	Yes 64
2	LD	F0	0	R1	6	10	11	Store1	No
2	MULTD	F4	F0	F2	7	15	16	Store2	No
2	SD	F4	0	R1	8	19	20	Store3	Yes 64 Multi1

Reservation Stations:

Time	Name	Busy	Op	Vj	Vk	Qj	Qk	Code:
Add1	No							LD F0 0 R1
Add2	No							MULTD F4 F0 F2
Add3	No							SD F4 0 R1
Mult1	Yes	Multd				R(F2)	Load3	SUBI R1 R1 #8
Mult2	No							BNEZ R1 Loop

Register result status

Clock	R1	F0	F2	F4	F6	F8	F10	F12	...	F30
20	56	Fu	Load1	Multi1						

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

105

Why can Tomasulo overlap iterations of loops?

- Register renaming
 - Multiple iterations use different physical destinations for registers (dynamic loop unrolling).
- Reservation stations
 - Permit instruction issue to advance past integer control flow operations
 - Also buffer old values of registers - totally avoiding the WAR stall that we saw in the scoreboard.
- Other perspective: Tomasulo building data flow dependency graph on the fly.

106

Major Advantages of Tomasulo's Scheme

- (1) The distribution of the hazard detection logic
 - distributed reservation stations and the Common Data Bus (CDB)
 - If multiple instructions waiting on single result, & each instruction has other operand, then instructions can be released simultaneously by broadcast on CDB
 - If a centralized register file were used, the units would have to read their results from the registers when register buses are available.
- (2) The elimination of stalls for WAW and WAR hazards

107

What about Precise Interrupts?

- State of machine looks as if no instruction beyond faulting instructions has issued
- Tomasulo had:
 - In-order issue, out-of-order execution, and out-of-order completion
- Need to "fix" the out-of-order completion aspect so that we can find precise breakpoint in instruction stream.

108

Relationship between Precise Interrupts and Speculation:

- **Speculation:** guess and check
- Important for branch prediction:
 - Need to “take our best shot” at predicting branch direction.
- If we speculate and are wrong, need to back up and restart execution to point at which we predicted incorrectly:
 - This is exactly same as precise exceptions!
- Technique for both precise interrupts/exceptions and speculation: **in-order completion or commit**

109

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

110

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

111

Reorder Buffer

- Issue:
 - Allocate **RS** (Reservation Stations) and **ReOrder Buffer (ROB)**, read available operands
- Execute:
 - Begin execution when operand values are available
- Write result:
 - Write result and **ROB tag on CDB** (Common Data Bus)
- Commit:
 - When **ROB** reaches head of **ROB**, update register
 - When a mispredicted branch reaches head of **ROB**, discard all entries

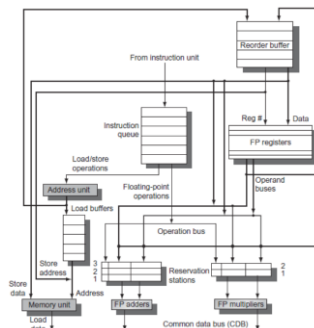
112

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

113

Reorder Buffer



114

Reorder Buffer

Reorder buffer					
Entry	Busy	Instruction	State	Destination	Value
1	No	f1d	F6, 32(x2)	Commit	F6
2	No	f1d	F2, 44(x3)	Commit	F2
3	Yes	fmul.d	F0, F2, F4	Write result	F0
4	Yes	fsub.d	F0, F2, F6	Write result	F0
5	Yes	fdiv.d	F0, F0, F6	Execute	F0
6	Yes	fadd.d	F6, F0, F2	Write result	F6

Reservation stations							
Name	Busy	Op	Vj	Vk	Qj	Qk	Dest A
Load1	No						
Load2	No						
Add1	No						
Add2	No						
Add3	No						
Mult1	No	fmul.d	Mem[44 + Regs[x3]]	Regs[*4]			#3
Mult2	Yes	fdiv.d		Mem[32 + Regs[x2]]		#3	#5

FP register status										
Field	#0	f1	f2	f3	f4	f5	f6	f7	f8	f10
Reorder #	3						6	4	5	
Busy	Yes	No	No	No	No	No	Yes	---	Yes	Yes

115

116

Speculation to greater ILP

- Greater ILP: Overcome control dependence by hardware speculating on outcome of branches and executing program as if guesses were correct
 - Speculation \Rightarrow fetch, issue, and execute instructions as if branch predictions were always correct
 - Dynamic scheduling \Rightarrow only fetches and issues instructions
- Essentially a data flow execution model: Operations execute as soon as their operands are available

Speculation to greater ILP

- 3 components of HW-based speculation:
 1. Dynamic branch prediction to choose which instructions to execute
 2. Speculation to allow execution of instructions before control dependences are resolved
 - + ability to undo effects of incorrectly speculated sequence
 3. Dynamic scheduling to deal with scheduling of different combinations of basic blocks

117

118

Adding Speculation to Tomasulo

- Must separate execution from allowing instruction to finish or “commit”
- This additional step called instruction commit
- When an instruction is no longer speculative, allow it to update the register file or memory
- Requires additional set of buffers to hold results of instructions that have finished execution but have not committed
- This reorder buffer (ROB) is also used to pass results among instructions that may be speculated

Reorder Buffer (ROB)

- In non-speculative Tomasulo’s algorithm, once an instruction writes its result, any subsequently issued instructions will find result in the register file
- With speculation, the register file is not updated until the instruction commits
 - (we know definitively that the instruction should execute)
- Thus, the ROB supplies operands in interval between completion of instruction execution and instruction commit
 - ROB is a source of operands for instructions, just as reservation stations (RS) provide operands in Tomasulo’s algorithm
 - ROB extends architected registers like RS

119

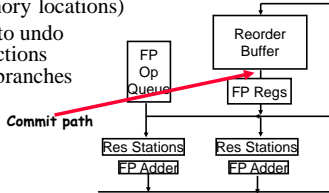
Reorder Buffer Entry Fields

- Each entry in the ROB contains four fields:
 1. Instruction type
 - a branch (has no destination result), a store (has a memory address destination), or a register operation (ALU operation or load, which has register destinations)
 2. Destination
 - Register number (for loads and ALU operations) or memory address (for stores) where the instruction result should be written
 3. Value
 - Value of instruction result until the instruction commits
 4. Ready
 - Indicates that instruction has completed execution, and the value is ready

120

Reorder Buffer Operation

- Holds instructions in **FIFO** order, exactly as issued
- When instructions complete, results placed into **ROB**
 - Supplies operands to other instruction between execution complete & commit \Rightarrow more registers like **RS**
 - Tag results with **ROB** buffer number instead of reservation station
- Instructions **commit** \Rightarrow values at head of **ROB** placed in registers (or memory locations)
- As a result, easy to undo speculated instructions on mispredicted branches or on exceptions

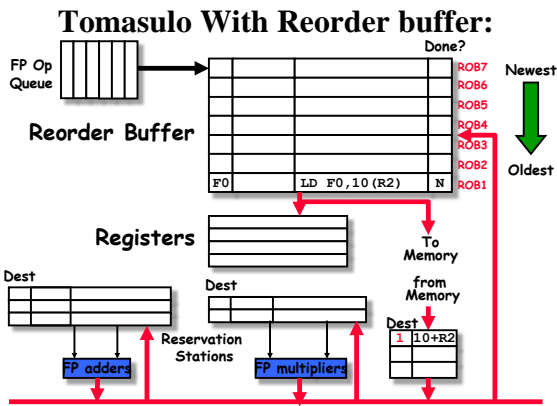


121

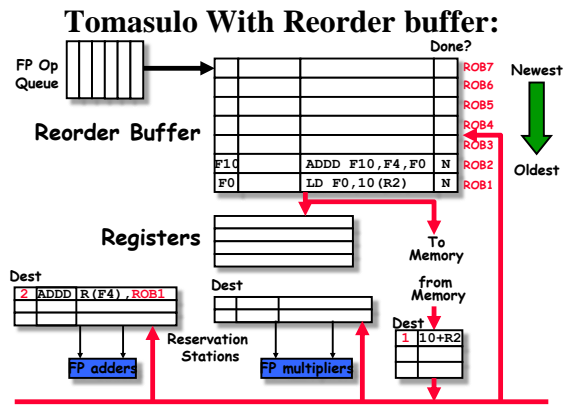
Recall: 4 Steps of Speculative Tomasulo Algorithm

1. **Issue**—get instruction from FP Op Queue
If reservation station and reorder buffer slot free, issue instr & send operands & reorder buffer no. for destination (this stage sometimes called “dispatch”)
2. **Execution**—operate on operands (EX)
When both operands ready then execute; if not ready, watch CDB for result; when both in reservation station, execute; checks RAW (sometimes called “issue”)
3. **Write result**—finish execution (WB)
Write on Common Data Bus to all awaiting FUs & reorder buffer; mark reservation station available.
4. **Commit**—update register with reorder result
When instr. at head of reorder buffer & result present, update register with result (or store to memory) and remove instr from reorder buffer. Mispredicted branch flushes reorder buffer (sometimes called “graduation”)

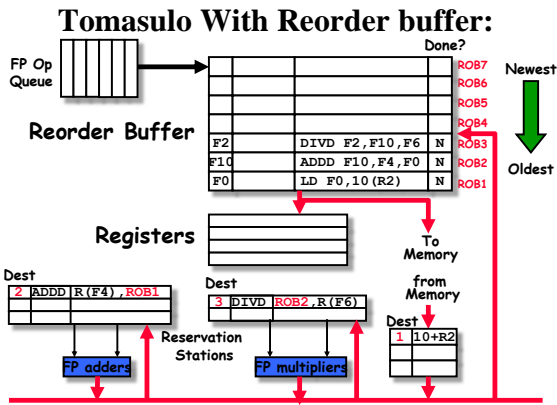
122



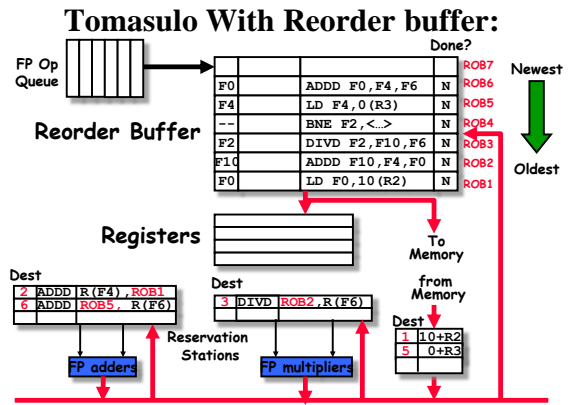
123



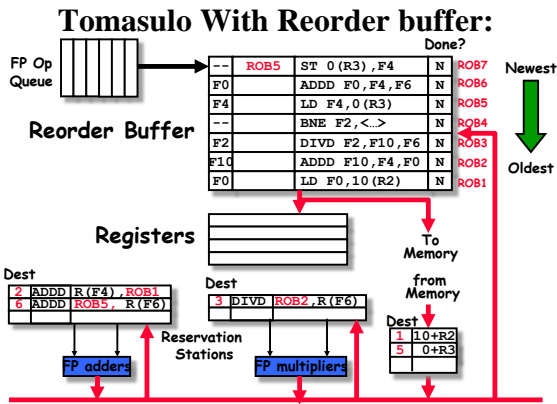
124



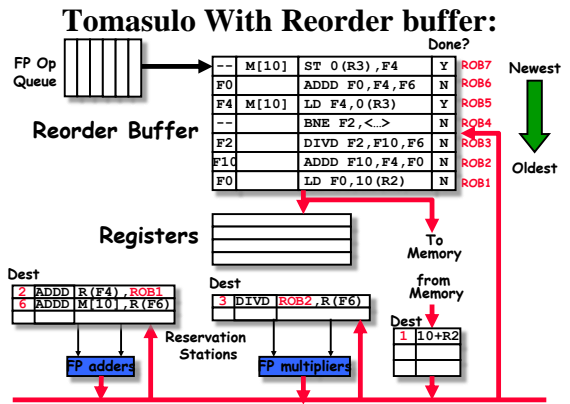
125



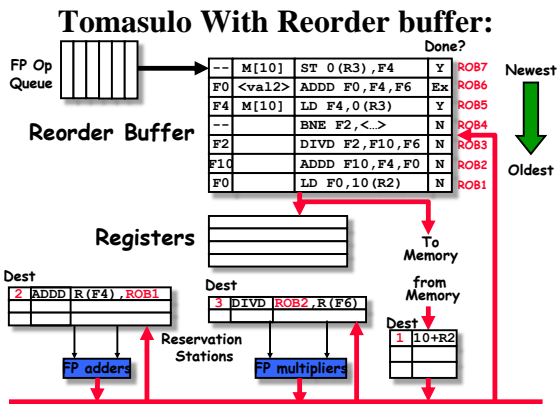
126



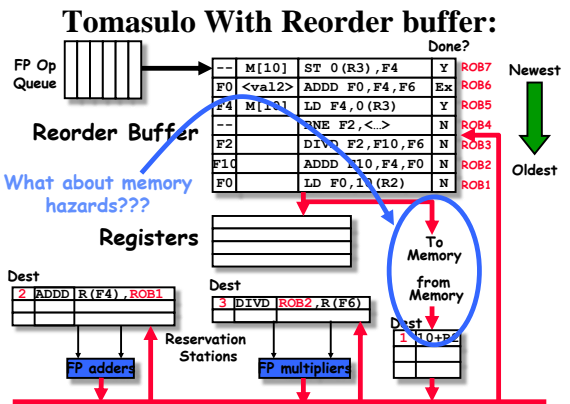
127



128



129



130

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 dependence 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

131

Exceptions and Interrupts

- IBM 360/91 invented “imprecise interrupts”
 - Computer stopped at this PC; its likely close to this address
 - Not so popular with programmers
 - Also, what about Virtual Memory? (Not in IBM 360)
- Technique for both precise interrupts/exceptions and speculation: in-order completion and in-order commit
 - If we speculate and are wrong, need to back up and restart execution to point at which we predicted incorrectly
 - This is exactly same as need to do with precise exceptions
- Exceptions are handled by not recognizing the exception until instruction that caused it is ready to commit in ROB
 - If a speculated instruction raises an exception, the exception is recorded in the ROB
 - This is why reorder buffers in all new processors

132

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 Cortex-A53
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

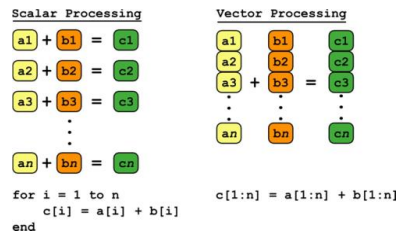
133

134

Getting CPI < 1: Issuing Multiple Instructions/Cycle

- Vector Processing:** Explicit coding of independent loops as operations on large vectors of numbers
 - Multimedia instructions being added to many processors
- Superscalar:** varying no. instructions/cycle (1 to 8), scheduled by compiler or by HW (Tomasulo)
 - IBM PowerPC, Sun UltraSparc, Pentium 4
- (Very) Long Instruction Words (V)LIW:** fixed number of instructions (4-16) scheduled by the compiler; put ops into wide templates
 - Intel Architecture-64 (IA-64) 64-bit address
 - Renamed: "Explicitly Parallel Instruction Computer (EPIC)"
- Anticipated success of multiple instructions lead to **Instructions Per Clock_{cycle} (IPC)** vs. CPI

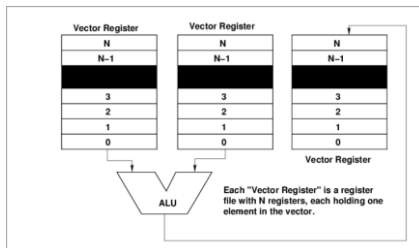
Vector Processor



135

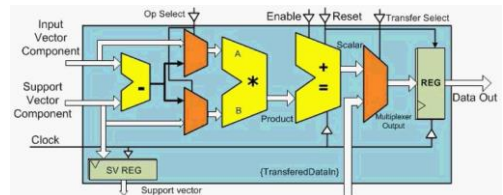
136

Vector Processor



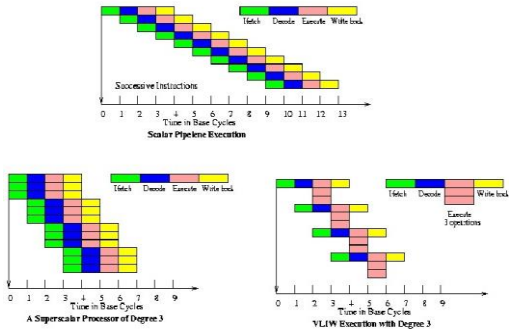
137

Vector Processor



138

Superscalar and VLIW



139

Vector Processor

Scalar Processing

$$a_1 + b_1 = c_1$$

$$a_2 + b_2 = c_2$$

$$a_3 + b_3 = c_3$$

$$\vdots$$

$$a_n + b_n = c_n$$

for $i = 1$ to n
 $c[i] = a[i] + b[i]$
end

Vector Processing

$$\begin{bmatrix} a_1 \\ a_2 \\ a_3 \\ \vdots \\ a_n \end{bmatrix} + \begin{bmatrix} b_1 \\ b_2 \\ b_3 \\ \vdots \\ b_n \end{bmatrix} = \begin{bmatrix} c_1 \\ c_2 \\ c_3 \\ \vdots \\ c_n \end{bmatrix}$$

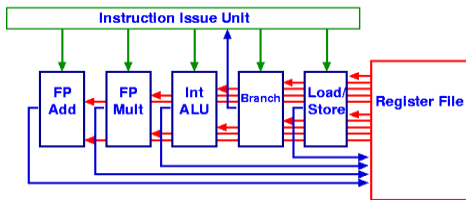
$c[1:n] = a[1:n] + b[1:n]$

140

VLIW (Very Long Instruction Word)

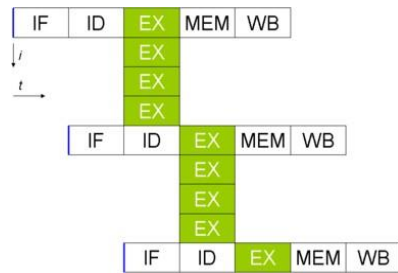
Instruction Format

FP Add FP Mult Int ALU Branch Load/Store



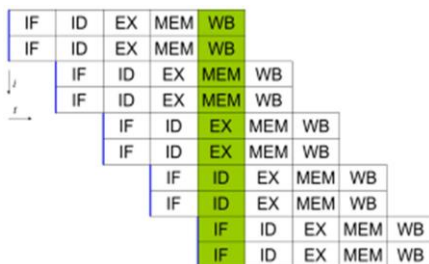
141

VLIW Pipeline



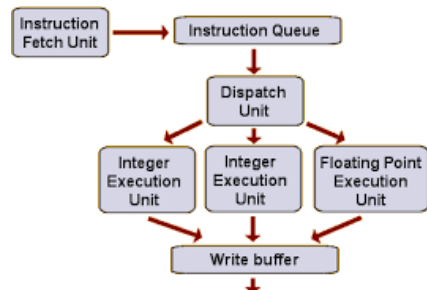
142

Superscalar Pipeline



143

Superscalar



144

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

145

VLIW Processors

Memory reference 1	Memory reference 2	FP operation 1	FP operation 2	Integer operation/branch
f1d f0,0(x1)	f1d f6,-8(x1)			
f1d f10,-16(x1)	f1d f14,-24(x1)			
f1d f18,-32(x1)	f1d f22,-40(x1)	fadd.d f4,f0,f2	fadd.d f8,f6,f2	
f1d f26,-48(x1)		fadd.d f12,f0,f2	fadd.d f16,f14,f2	
		fadd.d f20,f18,f2	fadd.d f24,f22,f2	
fsd f4,0(x1)	fsd f8,-8(x1)	fadd.d f28,f26,f24		
fsd f12,-16(x1)	fsd f16,-24(x1)			add! x1,x1,-56
fsd f20,24(x1)	fsd f24,16(x1)			
fsd f28,8(x1)				bne x1,x2,loop

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

146

VLIW: Very Large Instruction Word

- Each “instruction” has explicit coding for multiple operations
 - In IA-64, grouping called a “packet”
 - In Transmeta, grouping called a “molecule” (with “atoms” as ops)
- Tradeoff instruction space for simple decoding
 - The long instruction word has room for many operations
 - By definition, all the operations the compiler puts in the long instruction word are independent => execute in parallel
 - E.g., 2 integer operations, 2 FP ops, 2 Memory refs, 1 branch
 - 16 to 24 bits per field => 7*16 or 112 bits to 7*24 or 168 bits wide
 - Need compiling technique that schedules across several branches

147

Recall: Unrolled Loop that Minimizes Stalls for Scalar

1	Loop: L.D	F0,0(R1)	L.D to ADD.D: 1 Cycle
2	L.D	F6,-8(R1)	ADD.D to S.D: 2 Cycles
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	

14 clock cycles, or 3.5 per iteration

148

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)

149

Problems with 1st Generation VLIW

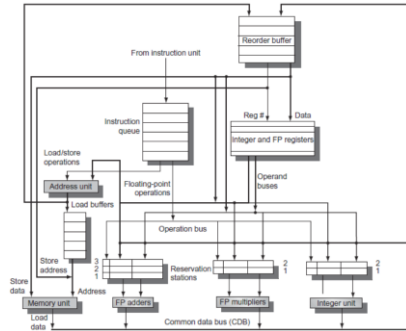
- Increase in code size
 - generating enough operations in a straight-line code fragment requires ambitiously unrolling loops
 - whenever VLIW instructions are not full, unused functional units translate to wasted bits in instruction encoding
- Operated in lock-step; no hazard detection HW
 - a stall in any functional unit pipeline caused entire processor to stall, since all functional units must be kept synchronized
 - Compiler might predict function units, but caches hard to predict
- Binary code compatibility
 - Pure VLIW => different numbers of functional units and unit latencies require different versions of the code

150

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
- Issue logic is the bottleneck in dynamically scheduled superscalars

Overview of Design



151

152

Multiple Issue

- 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
- To simplify RS allocation:
 - 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

153

154

Example

```

Loop: ld x2,0(x1)      //x2=array element
      addi x2,x2,1    //increment x2
      sd x2,0(x1)     //store result
      addi x1,x1,8    //increment pointer
      bne x2,x3,Loop //branch if not last
    
```

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 x2,0(x1)	1	2	3	4	First issue
1	addi x2,x2,1	1	5	6	7	Wait for ld
1	sd x2,0(x1)	2	3	7	8	Wait for addi
1	addi x1,x1,8	2	3	4	5	Execute directly
1	bne x2,x3,Loop	3	7	8	9	Wait for addi
2	ld x2,0(x1)	4	8	9	10	Wait for bne
2	addi x2,x2,1	4	11	12	13	Wait for ld
2	sd x2,0(x1)	5	9	13	14	Wait for addi
2	addi x1,x1,8	5	8	9	10	Wait for bne
2	bne x2,x3,Loop	6	13	14	15	Wait for addi
3	ld x2,0(x1)	7	14	15	16	Wait for bne
3	addi x2,x2,1	7	17	18	19	Wait for ld
3	sd x2,0(x1)	8	15	19	20	Wait for addi
3	addi x1,x1,8	8	14	15	16	Wait for bne
3	bne x2,x3,Loop	9	19	20	21	Wait for addi

155

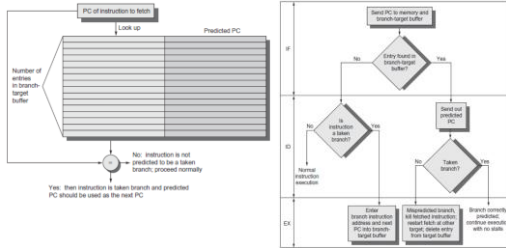
Example (Multiple Issue with Speculation)

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 x2,0(x1)	1	2	3	4	5	First issue
1	addi x2,x2,1	1	5	6	7	7	Wait for ld
1	sd x2,0(x1)	2	3	7	8	8	Wait for addi
1	addi x1,x1,8	2	3	4	5	5	Commit in order
1	bne x2,x3,Loop	3	7	8	9	9	Wait for addi
2	ld x2,0(x1)	4	5	6	7	7	No execute delay
2	addi x2,x2,1	4	8	9	10	10	Wait for ld
2	sd x2,0(x1)	5	6	10	11	11	Wait for addi
2	addi x1,x1,8	5	6	7	8	8	Commit in order
2	bne x2,x3,Loop	6	10	11	12	12	Wait for addi
3	ld x2,0(x1)	7	8	9	10	10	Earliest possible
3	addi x2,x2,1	7	11	12	13	13	Wait for ld
3	sd x2,0(x1)	8	9	9	10	10	Wait for addi
3	addi x1,x1,8	8	9	10	11	11	Executes earlier
3	bne x2,x3,Loop	9	13	14	15	15	Wait for addi

156

Branch-Target Buffer

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



157

Branch Folding

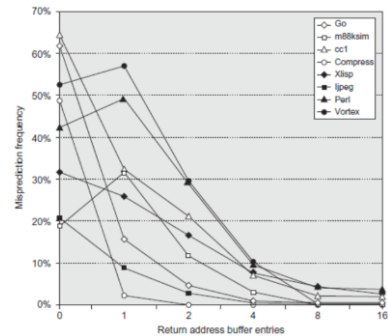
- Optimization:
 - Larger branch-target buffer
 - Add target instruction into buffer to deal with longer decoding time required by larger buffer
 - “Branch folding”

158

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

Return Address Predictor



159

160

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

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
 - Simple approach: deallocate virtual register when next instruction writes to its mapped architecturally-visible register

161

162

Integrated Issue and Renaming

- Combining instruction issue with register renaming:
 - Issue logic pre-reserves enough physical registers for the bundle
 - 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

Instr. #	Instruction	Physical register assigned or destination	Instruction with physical register numbers	Rename map changes
1	add x1, x2, x3	p32	add p32, p2, p3	x1-> p32
2	sub x1, x1, x2	p33	sub p33, p32, p2	x1->p33
3	add x2, x1, x2	p34	add p34, p33, x2	x2->p34
4	sub x1, x3, x2	p35	sub p35, p3, p34	x1->p35
5	add x1, x1, x2	p36	add p36, p35, p34	x1->p36
6	sub x1, x3, x1	p37	sub p37, p3, p36	x1->p37

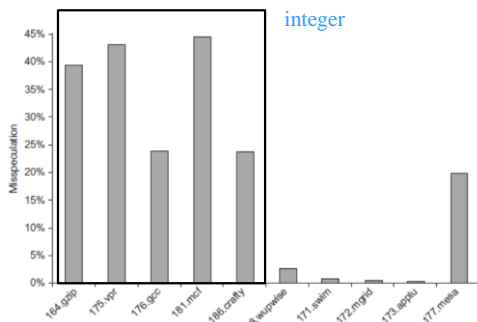
163

How Much Speculation?

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

164

How Much Speculation?



165

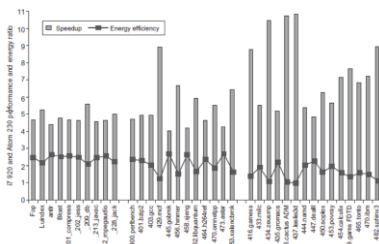
Energy Efficiency

- Value prediction
 - Uses:
 - Loads that load from a constant pool
 - Instruction that produces a value from a small set of values
 - Not incorporated into modern processors
 - Similar idea--*address aliasing prediction*--is used on some processors to determine if two stores or a load and a store reference the same address to allow for reordering

166

Fallacies and Pitfalls

- It is easy to predict the performance/energy efficiency of two different versions of the same ISA if we hold the technology constant



167

Fallacies and Pitfalls

- Processors with lower CPIs / faster clock rates will also be faster

Processor	Implementation technology	Clock rate	Power	SPECint2006 base	SPECint2006 baseline
Intel Pentium 4 670	90 nm	3.8 GHz	115 W	11.5	12.2
Intel Itanium 2	90 nm	1.66 GHz	104 W approx. 70 W one core	14.5	17.3
Intel i7 920	45 nm	3.3 GHz	130 W total approx. 80 W one core	35.5	38.4

- Pentium 4 had higher clock, lower CPI
- Itanium had same CPI, lower clock

168

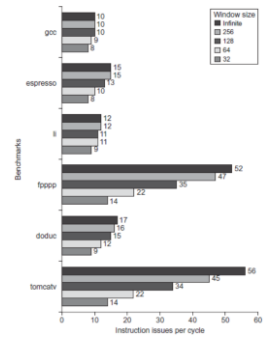
Fallacies and Pitfalls

- Sometimes bigger and dumber is better
 - Pentium 4 and Itanium were advanced designs, but could not achieve their peak instruction throughput because of relatively small caches as compared to i7
- And sometimes smarter is better than bigger and dumber
 - TAGE branch predictor outperforms gshare with less stored predictions

169

Fallacies and Pitfalls

- Believing that there are large amounts of ILP available, if only we had the right techniques



170