

CS152  
Computer Architecture and Engineering  
Lecture 16

Dynamic Scheduling, Speculation, and ILP

March 31, 1999

John Kubiawicz (<http://cs.berkeley.edu/~kubitron>)

lecture slides: <http://www-inst.eecs.berkeley.edu/~cs152/>

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.1

Review: Compiler techniques for parallelism

- Loop unrolling  $\Rightarrow$  Multiple iterations of loop in software:
  - Amortizes loop overhead over several iterations
  - Gives more opportunity for scheduling around stalls
- Software Pipelining  $\Rightarrow$  Take one instruction from each of several iterations of the loop
  - Software overlapping of loop iterations
  - Today will show hardware overlapping of loop iterations
- Very Long Instruction Word machines (VLIW)  $\Rightarrow$  Multiple operations coded in single, long instruction
  - Requires sophisticated compiler to decide which operations can be done in parallel
  - Trace scheduling  $\Rightarrow$  find common path and schedule code as if branches didn't exist (+ add "fixup code")
- **All of these require additional registers**

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.2

Review: Dynamic hardware for out-of-order execution

- HW exploitation of ILP
  - Works when can't know dependence at compile time.
  - Code for one machine runs well on another
- Scoreboard (ala CDC 6600 in 1963)
  - Centralized control structure
  - No register renaming, no forwarding
  - Pipeline stalls for WAR and WAW hazards.
  - **Are these fundamental limitations???** (No)
- Reservation stations (ala IBM 360/91 in 1966)
  - Distributed control structures
  - **Implicit** renaming of registers (dispatched pointers)
  - WAR and WAW hazards **eliminated** by register renaming
  - Results broadcast to all reservation stations for RAW
  - **Reservation stations are like additional registers**

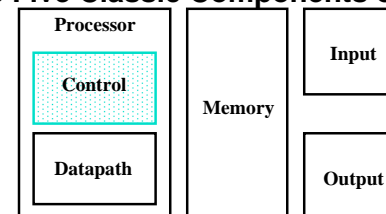
3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.3

The Big Picture: Where are We Now?

- The Five Classic Components of a Computer



- Today's Topics:
  - Recap last lecture
  - Hardware loop unrolling with Tomasulo algorithm
  - Administrivia
  - Speculation, branch prediction
  - Reorder buffers

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.4

## Three Stages of Tomasulo Algorithm

### 1. Issue—get instruction from FP Op Queue

If reservation station free (no structural hazard), control issues instr & sends operands (renames registers).

### 2. Execution—operate on operands (EX)

When both operands ready then execute; if not ready, watch Common Data Bus for result

### 3. Write result—finish execution (WB)

Write on Common Data Bus to all awaiting units; mark reservation station available

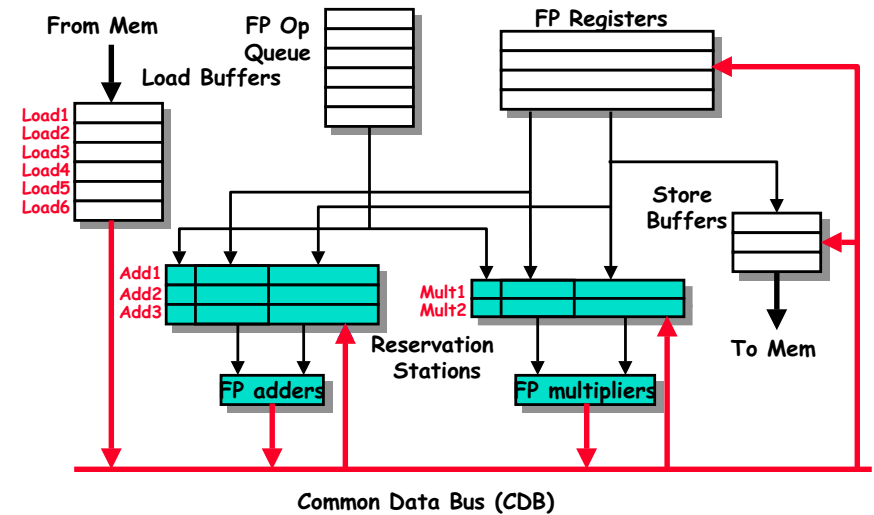
- Normal data bus: data + destination (“go to” bus)
- Common data bus: data + source (“come from” bus)
  - 64 bits of data + 4 bits of Functional Unit source address
  - Write if matches expected Functional Unit (produces result)
  - Does the broadcast

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.5

## Tomasulo Organization



3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.6

## Tomasulo Loop Example

```

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

- Assume Multiply takes 4 clocks
- Assume first load takes 8 clocks (cache miss), second load takes 1 clock (hit)
- To be clear, will show clocks for SUBI, BNEZ
- Reality: integer instructions ahead

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.7

## Loop Example

Instruction status:

ITER	Instruction	j	k	Exec Write			Busy	Addr	Fu
				Issue	Comp	Result			
1	LD	F0	0	R1			No		
1	MULTD	F4	F0	F2			No		
1	SD	F4	0	R1			No		
2	LD	F0	0	R1			No		
2	MULTD	F4	F0	F2			No		
2	SD	F4	0	R1			No		

Reservation Stations:

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

Register result status

Clock	R1	F0	F2	F4	F6	F8	F10	F12	...	F30
0	80	Fu								

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.8

## Loop Example Cycle 1

Instruction status:

ITER	Instruction	j	k	Issue	Comp	Result	Busy	Addr	Fu
1	LD F0 0 R1			1			Load1	Yes 80	
1	MULTD F4 F0 F2						Load2	No	
1	SD F4 0 R1						Load3	No	
2	LD F0 0 R1						Store1	No	
2	MULTD F4 F0 F2						Store2	No	
2	SD F4 0 R1						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
Mult1	No							SUBI R1 R1 #8
Mult2	No							BNEZ R1 Loop

Register result status

Clock	R1	F0	F2	F4	F6	F8	F10	F12	...	F30
1	80	Fu	Load1							

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.9

## Loop Example Cycle 2

Instruction status:

ITER	Instruction	j	k	Issue	Comp	Result	Busy	Addr	Fu
1	LD F0 0 R1			1			Load1	Yes 80	
1	MULTD F4 F0 F2						Load2	No	
1	SD F4 0 R1						Load3	No	
2	LD F0 0 R1						Store1	No	
2	MULTD F4 F0 F2						Store2	No	
2	SD F4 0 R1						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
Mult1	Yes	Multd				R(F2)	Load1	SUBI R1 R1 #8
Mult2	No							BNEZ R1 Loop

Register result status

Clock	R1	F0	F2	F4	F6	F8	F10	F12	...	F30
2	80	Fu	Load1	Mult1						

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.10

## Loop Example Cycle 3

Instruction status:

ITER	Instruction	j	k	Issue	Comp	Result	Busy	Addr	Fu
1	LD F0 0 R1			1			Load1	Yes 80	
1	MULTD F4 F0 F2						Load2	No	
1	SD F4 0 R1						Load3	No	
2	LD F0 0 R1						Store1	Yes 80	Mult1
2	MULTD F4 F0 F2						Store2	No	
2	SD F4 0 R1						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
Mult1	Yes	Multd				R(F2)	Load1	SUBI R1 R1 #8
Mult2	No							BNEZ R1 Loop

Register result status

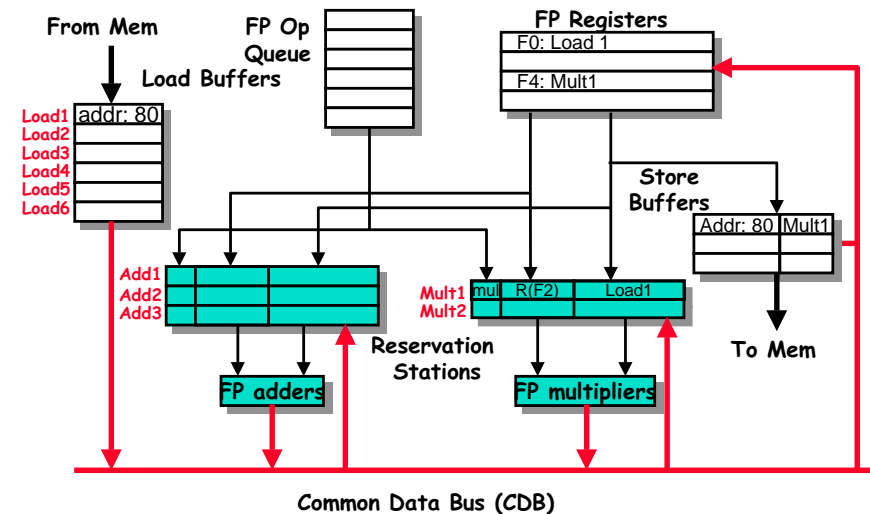
Clock	R1	F0	F2	F4	F6	F8	F10	F12	...	F30
3	80	Fu	Load1	Mult1						

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.11

## What does this mean physically?



3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.12

° Implicit renaming sets up "DataFlow" graph

## Loop Example Cycle 4

### Instruction status:

		Exec Write						
ITER	Instruction	j	k	Issue	CompResult	Busy	Addr	Fu
1	LD F0	0	R1	1		Load1	Yes 80	
1	MULTD F4	F0	F2	2		Load2	No	
1	SD F4	0	R1	3		Load3	No	
2	LD F0	0	R1			Store1	Yes 80	Mult1
2	MULTD F4	F0	F2			Store2	No	
2	SD F4	0	R1			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
Mult1	Yes	Multid				R(F2)	Load1	SUBI R1 R1 #8
Mult2	No							BNEZ R1 Loop

### Register result status

Clock	R1	F0	F2	F4	F6	F8	F10	F12	...	F30
4	80	Fu	Load1	Mult1						

### ° Dispatching SUBI Instruction

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.13

## Loop Example Cycle 5

### Instruction status:

		Exec Write						
ITER	Instruction	j	k	Issue	CompResult	Busy	Addr	Fu
1	LD F0	0	R1	1		Load1	Yes 80	
1	MULTD F4	F0	F2	2		Load2	No	
1	SD F4	0	R1	3		Load3	No	
2	LD F0	0	R1			Store1	Yes 80	Mult1
2	MULTD F4	F0	F2			Store2	No	
2	SD F4	0	R1			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
Mult1	Yes	Multid				R(F2)	Load1	SUBI R1 R1 #8
Mult2	No							BNEZ R1 Loop

### Register result status

Clock	R1	F0	F2	F4	F6	F8	F10	F12	...	F30
5	72	Fu	Load1	Mult1						

### ° And, BNEZ instruction

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.14

## Loop Example Cycle 6

### Instruction status:

		Exec Write						
ITER	Instruction	j	k	Issue	CompResult	Busy	Addr	Fu
1	LD F0	0	R1	1		Load1	Yes 80	
1	MULTD F4	F0	F2	2		Load2	Yes 72	
1	SD F4	0	R1	3		Load3	No	
2	LD F0	0	R1	6		Store1	Yes 80	Mult1
2	MULTD F4	F0	F2			Store2	No	
2	SD F4	0	R1			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
Mult1	Yes	Multid				R(F2)	Load1	SUBI R1 R1 #8
Mult2	No							BNEZ R1 Loop

### Register result status

Clock	R1	F0	F2	F4	F6	F8	F10	F12	...	F30
6	72	Fu	Load2	Mult1						

### ° Notice that F0 never sees Load from location 80

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.15

## Loop Example Cycle 7

### Instruction status:

		Exec Write						
ITER	Instruction	j	k	Issue	CompResult	Busy	Addr	Fu
1	LD F0	0	R1	1		Load1	Yes 80	
1	MULTD F4	F0	F2	2		Load2	Yes 72	
1	SD F4	0	R1	3		Load3	No	
2	LD F0	0	R1	6		Store1	Yes 80	Mult1
2	MULTD F4	F0	F2	7		Store2	No	
2	SD F4	0	R1			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
Mult1	Yes	Multid				R(F2)	Load1	SUBI R1 R1 #8
Mult2	Yes	Multid				R(F2)	Load2	BNEZ R1 Loop

### Register result status

Clock	R1	F0	F2	F4	F6	F8	F10	F12	...	F30
7	72	Fu	Load2	Mult2						

### ° Register file completely detached from iteration 1

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.14

## Loop Example Cycle 8

### Instruction status:

ITER	Instruction	j	k	Exec Write		Issue	CompResult	Busy	Addr	Fu
				Op	Rs					
1	LD	F0	0	R1	1			Load1	Yes	80
1	MULTD	F4	F0	F2	2			Load2	Yes	72
1	SD	F4	0	R1	3			Load3	No	
2	LD	F0	0	R1	6			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:	SI S2 RS		
									Op	Rs	Op
Add1	No							LD	F0	0	R1
Add2	No							MULTD	F4	F0	F2
Add3	No							SD	F4	0	R1
Mult1	Yes	Multd			R(F2)	Load1		SUBI	R1	R1	#8
Mult2	Yes	Multd			R(F2)	Load2		BNEZ	R1	Loop	

### Register result status

Clock	R1	F0	F2	F4	F6	F8	F10	F12	...	F30
8	72	Fu	Load2	Mult2						

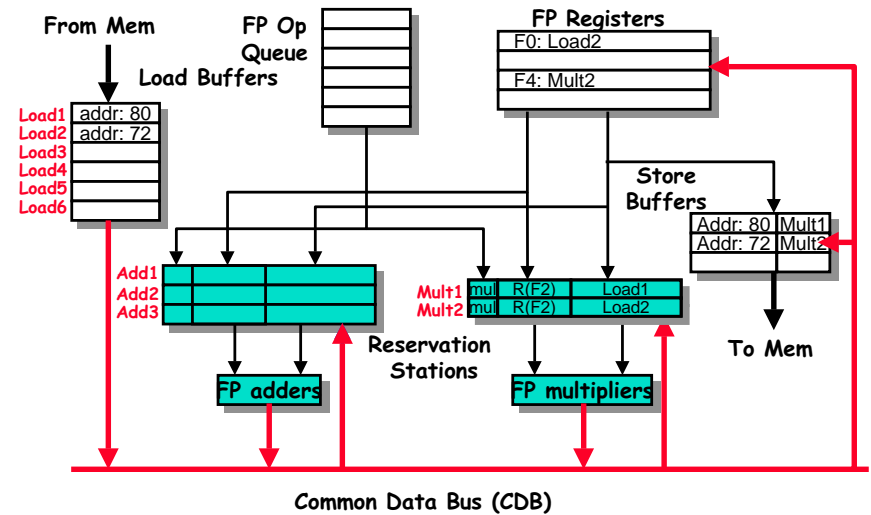
◦ First and Second iteration completely overlapped

3/31/99

©UCB Spring 1999

CS152 / Kubiatiowicz  
Lec16.17

## What does this mean physically?



3/31/99

©UCB Spring 1999

CS152 / Kubiatiowicz  
Lec16.18

## Loop Example Cycle 9

### Instruction status:

ITER	Instruction	j	k	Exec Write		Issue	CompResult	Busy	Addr	Fu
				Op	Rs					
1	LD	F0	0	R1	1	9		Load1	Yes	80
1	MULTD	F4	F0	F2	2			Load2	Yes	72
1	SD	F4	0	R1	3			Load3	No	
2	LD	F0	0	R1	6			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:	SI S2 RS		
									Op	Rs	Op
Add1	No							LD	F0	0	R1
Add2	No							MULTD	F4	F0	F2
Add3	No							SD	F4	0	R1
Mult1	Yes	Multd			R(F2)	Load1		SUBI	R1	R1	#8
Mult2	Yes	Multd			R(F2)	Load2		BNEZ	R1	Loop	

### Register result status

Clock	R1	F0	F2	F4	F6	F8	F10	F12	...	F30
9	72	Fu	Load2	Mult2						

◦ Load1 completing: who is waiting?

Note: Dispatching SUBI

3/31/99

©UCB Spring 1999

CS152 / Kubiatiowicz  
Lec16.19

## Loop Example Cycle 10

### Instruction status:

ITER	Instruction	j	k	Exec Write		Issue	CompResult	Busy	Addr	Fu
				Op	Rs					
1	LD	F0	0	R1	1	9	10	Load1	No	
1	MULTD	F4	F0	F2	2			Load2	Yes	72
1	SD	F4	0	R1	3			Load3	No	
2	LD	F0	0	R1	6	10		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:	SI S2 RS		
									Op	Rs	Op
Add1	No							LD	F0	0	R1
Add2	No							MULTD	F4	F0	F2
Add3	No							SD	F4	0	R1
4	Mult1	Yes	Multd	M[80]	R(F2)			SUBI	R1	R1	#8
Mult2	Yes	Multd			R(F2)	Load2		BNEZ	R1	Loop	

### Register result status

Clock	R1	F0	F2	F4	F6	F8	F10	F12	...	F30
10	64	Fu	Load2	Mult2						

◦ Load2 completing: who is waiting?

Note: Dispatching BNEZ

3/31/99

©UCB Spring 1999

CS152 / Kubiatiowicz  
Lec16.20

## Loop Example Cycle 11

### Instruction status:

ITER	Instruction	j	k	Exec Write			Busy	Addr	Fu
				Issue	Comp	Result			
1	LD F0 0 R1			1	9	10	Load1	No	
1	MULTD F4 F0 F2			2			Load2	No	
1	SD F4 0 R1			3			Load3	Yes	64
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:	SI			S2			RS		
	Add1	No						LD	F0	0	R1						
	Add2	No						MULTD	F4	F0	F2						
	Add3	No						SD	F4	0	R1						
3	Mult1	Yes	Multd	M[80]	R(F2)			SUBI	R1	R1	#8						
4	Mult2	Yes	Multd	M[72]	R(F2)			BNEZ	R1	Loop							

### Register result status

Clock	R1	F0	F2	F4	F6	F8	F10	F12	...	F30
11	64	Fu	Load3	Mult2						

### ◦ Next load in sequence

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.21

## Loop Example Cycle 12

### Instruction status:

ITER	Instruction	j	k	Exec Write			Busy	Addr	Fu
				Issue	Comp	Result			
1	LD F0 0 R1			1	9	10	Load1	No	
1	MULTD F4 F0 F2			2			Load2	No	
1	SD F4 0 R1			3			Load3	Yes	64
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:	SI			S2			RS		
	Add1	No						LD	F0	0	R1						
	Add2	No						MULTD	F4	F0	F2						
	Add3	No						SD	F4	0	R1						
2	Mult1	Yes	Multd	M[80]	R(F2)			SUBI	R1	R1	#8						
3	Mult2	Yes	Multd	M[72]	R(F2)			BNEZ	R1	Loop							

### Reservation Stations:

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

### ◦ Why not issue third multiply?

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.22

## Loop Example Cycle 13

### Instruction status:

ITER	Instruction	j	k	Exec Write			Busy	Addr	Fu
				Issue	Comp	Result			
1	LD F0 0 R1			1	9	10	Load1	No	
1	MULTD F4 F0 F2			2			Load2	No	
1	SD F4 0 R1			3			Load3	Yes	64
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:	SI			S2			RS		
	Add1	No						LD	F0	0	R1						
	Add2	No						MULTD	F4	F0	F2						
	Add3	No						SD	F4	0	R1						
1	Mult1	Yes	Multd	M[80]	R(F2)			SUBI	R1	R1	#8						
2	Mult2	Yes	Multd	M[72]	R(F2)			BNEZ	R1	Loop							

### Reservation Stations:

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

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.23

## Loop Example Cycle 14

### Instruction status:

ITER	Instruction	j	k	Exec Write			Busy	Addr	Fu
				Issue	Comp	Result			
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			Load3	Yes	64
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:	SI			S2			RS		
	Add1	No						LD	F0	0	R1						
	Add2	No						MULTD	F4	F0	F2						
	Add3	No						SD	F4	0	R1						
0	Mult1	Yes	Multd	M[80]	R(F2)			SUBI	R1	R1	#8						
1	Mult2	Yes	Multd	M[72]	R(F2)			BNEZ	R1	Loop							

### Reservation Stations:

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

### ◦ Mult1 completing. Who is waiting?

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.24

## Loop Example Cycle 15

### Instruction status:

ITER	Instruction	j	k	Exec Write			Busy	Addr	Fu
				Issue	Comp	Result			
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			Load3	Yes 64
2	LD	F0	0	R1	6	10	11	Store1	Yes 80 [80]*R2
2	MULTD	F4	F0	F2	7	15		Store2	Yes 72 Mult2
2	SD	F4	0	R1	8			Store3	No

### Reservation Stations:

Time	Name	Busy	Op	Vj	Vk	Qj	Qk	Code:	SI	S2	RS
	Add1	No						LD	F0	0	R1
	Add2	No						MULTD	F4	F0	F2
	Add3	No						SD	F4	0	R1
	Mult1	No						SUBI	R1	R1	#8
0	Mult2	Yes	Multd	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?

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.25

## Loop Example Cycle 16

### Instruction status:

ITER	Instruction	j	k	Exec Write			Busy	Addr	Fu
				Issue	Comp	Result			
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			Load3	Yes 64
2	LD	F0	0	R1	6	10	11	Store1	Yes 80 [80]*R2
2	MULTD	F4	F0	F2	7	15	16	Store2	Yes 72 [72]*R2
2	SD	F4	0	R1	8			Store3	No

### Reservation Stations:

Time	Name	Busy	Op	Vj	Vk	Qj	Qk	Code:	SI	S2	RS
	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
16	64	Fu	Load3	Mult1						

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.26

## Loop Example Cycle 17

### Instruction status:

ITER	Instruction	j	k	Exec Write			Busy	Addr	Fu
				Issue	Comp	Result			
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			Load3	Yes 64
2	LD	F0	0	R1	6	10	11	Store1	Yes 80 [80]*R2
2	MULTD	F4	F0	F2	7	15	16	Store2	Yes 72 [72]*R2
2	SD	F4	0	R1	8			Store3	Yes 64 Mult1

### Reservation Stations:

Time	Name	Busy	Op	Vj	Vk	Qj	Qk	Code:	SI	S2	RS
	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
17	64	Fu	Load3	Mult1						

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.27

## Loop Example Cycle 18

### Instruction status:

ITER	Instruction	j	k	Exec Write			Busy	Addr	Fu
				Issue	Comp	Result			
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		Load3	Yes 64
2	LD	F0	0	R1	6	10	11	Store1	Yes 80 [80]*R2
2	MULTD	F4	F0	F2	7	15	16	Store2	Yes 72 [72]*R2
2	SD	F4	0	R1	8			Store3	Yes 64 Mult1

### Reservation Stations:

Time	Name	Busy	Op	Vj	Vk	Qj	Qk	Code:	SI	S2	RS
	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	Mult1						

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.28

## Loop Example Cycle 19

### Instruction status:

ITER	Instruction	j	k	Exec Write			Busy	Addr	Fu		
				Issue	Comp	Result					
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]*R2
2	SD	F4	0	R1	8	19		Store3	Yes	64	Mult1

### Reservation Stations:

Time	Name	Busy	Op	Vj	Vk	Qj	Qk	SI S2 RS			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	64	Fu	Load3	Mult1						

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.29

## Loop Example Cycle 20

### Instruction status:

ITER	Instruction	j	k	Exec Write			Busy	Addr	Fu		
				Issue	Comp	Result					
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	No		
2	SD	F4	0	R1	8	19	20	Store3	Yes	64	Mult1

### Reservation Stations:

Time	Name	Busy	Op	Vj	Vk	Qj	Qk	SI S2 RS			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	64	Fu	Load3	Mult1						

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.30

## Why can Tomasulo overlap iterations of loops?

### Register renaming

- Multiple iterations use different physical destinations for registers (dynamic loop unrolling).
- Replace static register names from code with dynamic register "pointers"
- Effectively increases size of register file
- Permit instruction issue to advance past integer control flow operations.

### Crucial: integer unit must "get ahead" of floating point unit so that we can issue multiple iterations

### Other idea: Tomasulo building "DataFlow" graph.

## Recall: Unrolled Loop That Minimizes Stalls

```

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

14 clock cycles, or 3.5 per iteration  
Used new registers => register renaming!

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.31

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.32

## Administrivia

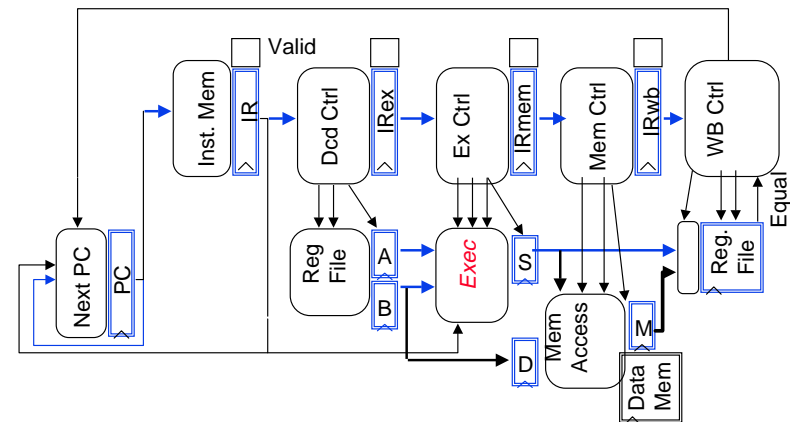
- Should be debugging Lab 5 by now!
- Remember: a *Working* processor is necessary for full credit...
- More info on some of the things that we have been talking about this week:
  - Computer Architecture: A Quantitative Approach by John Hennessy and David Patterson
- Next week: Memory systems
  - Start reading Chapter 7 (of your text) now...
  - Lab 6 will be using memory systems.

3/31/99

©UCB Spring 1999

CS152 / Kubiatiowicz  
Lec16.33

## Administrivia: Be careful about clock edges in lab5!



3/31/99

©UCB Spring 1999

CS152 / Kubiatiowicz  
Lec16.34

## Why issue in order?

- In-order issue permits us to analyze data flow of program
- This way, we know exactly which results should flow to which **subsequent** instructions
  - If we issued out-of-order, we would confuse RAW and WAR hazards!
  - The most advanced machines that I know of all issue in order.
- This idea works perfectly well “in principle” with multiple instructions issued per clock:
  - Need to multi-port “rename table” and be able to rename a sequence of instructions together
  - Need to be able to issue to multiple reservation stations in a single cycle.
  - Need to have 2x number of read ports and 1x number of write ports in register file.
- In-order issue can be serious bottleneck when issuing multiple instructions per clock-cycle

3/31/99

©UCB Spring 1999

CS152 / Kubiatiowicz  
Lec16.35

## Branches must be resolved quickly for loop overlap!

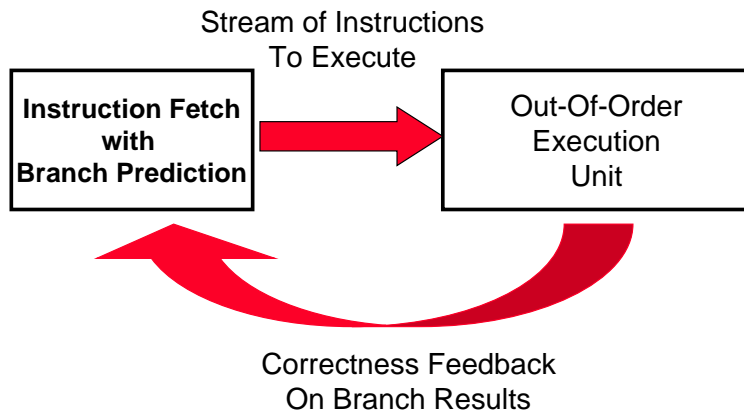
- In our example, we relied on the fact that branches were under control of “fast” integer unit in order to get overlap!
- ```
Loop:  LD      F0  0    R1
      MULTD  F4  F0  F2
      SD     F4  0    R1
      SUBI   R1  R1  #8
      BNEZ  R1   Loop
```
- What happens if branch depends on result of multd??
    - We completely lose all of our advantages!
    - Need to be able to “predict” branch outcome.
    - If we were to predict that branch was taken, this would be right most of the time.
  - Problem **much** worse for superscalar machines!

3/31/99

©UCB Spring 1999

CS152 / Kubiatiowicz  
Lec16.36

## Independent “Fetch” unit



- Instruction fetch decoupled from execution
- Often issue logic (+ rename) included with Fetch

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.37

## Prediction: Branches, Dependencies, Data

- Prediction has become essential to getting good performance from scalar instruction streams.
- We will discuss predicting branches. However, architects are not predicting data dependencies, actual data, and results of groups of instructions:
  - At what point does computation become a probabilistic operation + verification?
  - We are pretty close with control hazards already...
- 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.
- Prediction ⇒ Compressible information streams?

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.38

## Dynamic Branch Prediction

- Prediction could be “Static” (at compile time) or “Dynamic” (at runtime)
  - For our example, if we were to statically predict “taken”, we would only be wrong once each pass through loop
- Is dynamic branch prediction better than static branch prediction?
  - Seems to be. Still some debate to this effect
  - Today, lots of hardware being devoted to dynamic branch predictors.

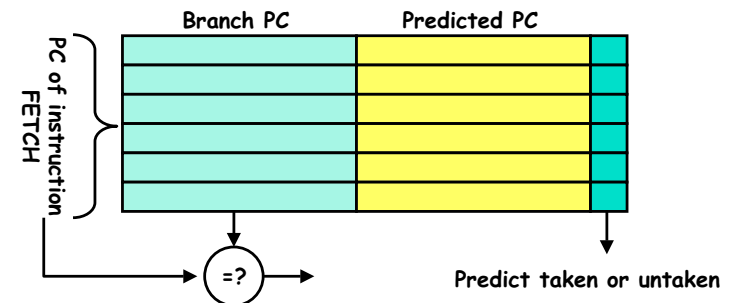
3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.39

## Simple dynamic prediction: Branch Target Buffer (BTB)

- Address of branch index to get prediction AND branch address (if taken)
  - Must check for branch match now, since can’t use wrong branch address
  - Grab predicted PC from table since may take several cycles to compute



- Update predicted PC when branch is actually resolved
- Return instruction addresses predicted with stack

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.40

## Dynamic Branch Prediction

- Performance =  $f(\text{accuracy, cost of misprediction})$
- 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

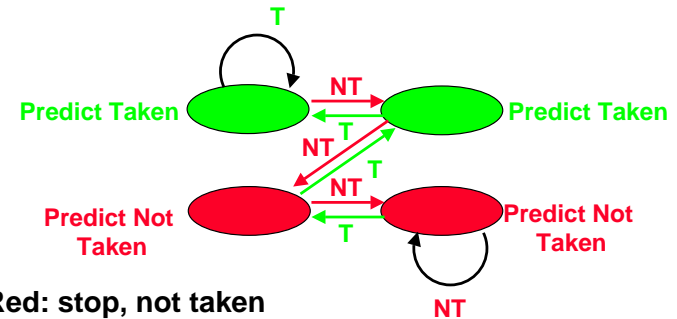
3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.41

## Dynamic Branch Prediction

- Solution: 2-bit scheme where change prediction only if get misprediction *twice*: (Figure 4.13, p. 264)



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

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.42

## BHT Accuracy

- Mispredict because either:
  - Wrong guess for that branch
  - Got branch history of wrong branch when index the table
- 4096 entry table programs vary from 1% misprediction (nasa7, tomcatv) to 18% (eqntott), with spice at 9% and gcc at 12%
- 4096 about as good as infinite table (in Alpha 211164)

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.43

## Correlating Branches

- Hypothesis: recent branches are correlated; that is, behavior of recently executed branches affects prediction of current branch
- Two possibilities; Current branch depends on:
  - Last  $m$  most recently executed branches anywhere in program  
Produces a "GA" (for "global address") in the Yeh and Patt classification (e.g. GAg)
  - Last  $m$  most recent outcomes of same branch.  
Produces a "PA" (for "per address") in same classification (e.g. PAG)
- Idea: record  $m$  most recently executed branches as taken or not taken, and use that pattern to select the proper branch history table entry
  - A single history table shared by all branches (appends a "g" at end, indexed by history value).
  - Address is used along with history to select table entry (appends a "p" at end of classification)
  - If only portion of address used, often appends an "s" to indicate "set-indexed" tables (i.e. GAs)

3/31/99

©UCB Spring 1999

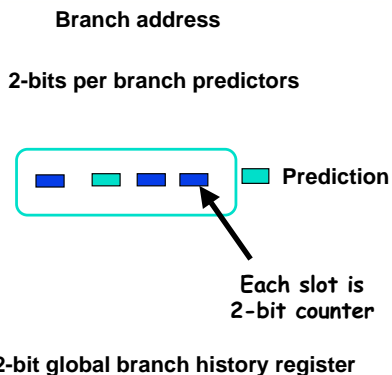
CS152 / Kubiawicz  
Lec16.44

## Correlating Branches

- For instance, consider global history, set-indexed BHT. That gives us a GAs history table.

### (2,2) GAs predictor

- First 2 means that we keep two bits of history
- Second means that we have 2 bit counters in each slot.
- Then behavior of recent branches selects between, say, four predictions of next branch, updating just that prediction
- Note that the original two-bit counter solution would be a (0,2) GAs predictor
- Note also that aliasing is possible here...

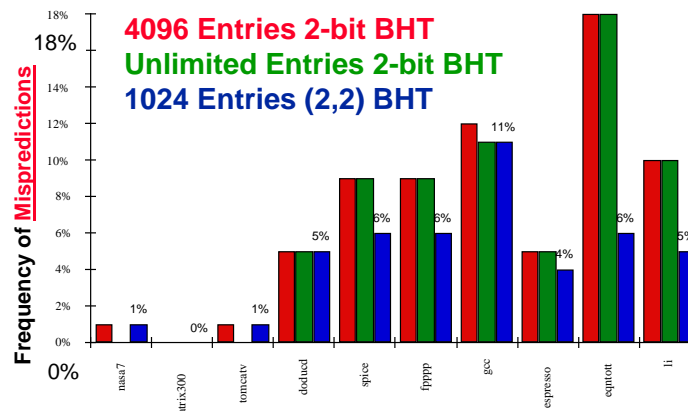


3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.45

## Accuracy of Different Schemes



3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.46

## HW support for More ILP

- Avoid branch prediction by turning branches into conditionally executed instructions:

**if (x) then A = B op C else NOP**

- If false, then neither store result nor cause exception
- Expanded ISA of Alpha, MIPS, PowerPC, SPARC have conditional move; PA-RISC can annul any following instr.
- EPIC: 64 1-bit condition fields selected so conditional execution

- Drawbacks to conditional instructions

- Still takes a clock even if “annulled”
- Stall if condition evaluated late
- Complex conditions reduce effectiveness; condition becomes known late in pipeline

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.47

## Now what about exceptions???

- Out-of-order commit really messes up our chance to get precise exceptions!
  - When committing results out-of-order, register file contains results from later instructions while earlier ones have not completed yet.
  - What if need to cause exception on one of those early instructions??
- Need to be able to “rollback” register file to consistent state
  - Remember that “precise” means that there is some PC such that: all instructions before have committed results, and none after have committed results.
- Big problem for branch prediction as well: What if prediction wrong??

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.48

## Relationship between precise interrupts and speculation:

- Speculation is a form of guessing.
- Important for branch prediction:
  - Need to “take our best shot” at predicting branch direction.
  - If we issue multiple instructions per cycle, lose lots of potential instructions otherwise:
    - Consider 4 instructions per cycle
    - If take single cycle to decide on branch, waste from 4 - 7 instruction slots!
- 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**

3/31/99

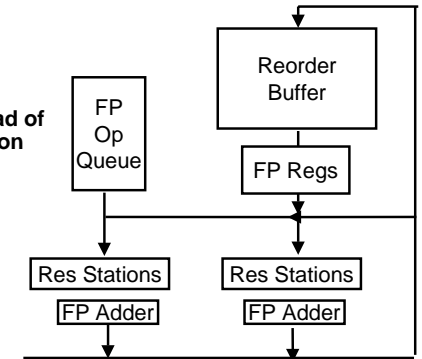
©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.49

## HW support for precise interrupts

- Need HW buffer for results of uncommitted instructions: **reorder buffer**

- 3 fields: instr, destination, value
- Reorder buffer can be operand source => more registers like RS
- Use reorder buffer number instead of reservation station when execution completes
- Supplies operands between execution complete & commit
- Once operand commits, result is put into register
- Instructions **commit**
- As a result, its easy to undo speculated instructions on mispredicted branches **or on exceptions**



3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.50

## Four Steps of Speculative Tomasulo Algorithm

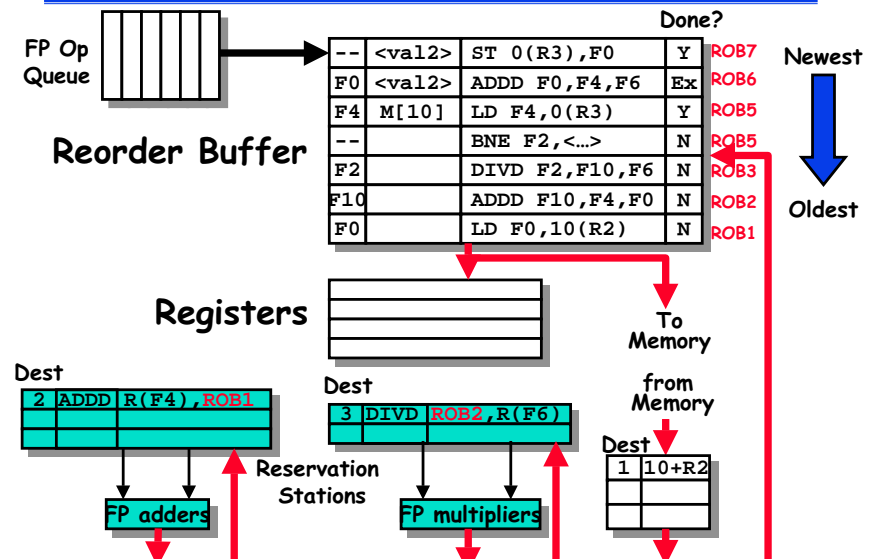
- 1. Issue**—get instruction from FP Op Queue
  - If reservation station **and reorder buffer slot** free, issue instr & send operands & **reorder buffer no. for destination** (this stage sometimes called “dispatch”)
- 2. Execution**—operate on operands (EX)
  - When both operands ready then execute; 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 or interrupt flushes reorder buffer (sometimes called “graduation”)

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.51

## Tomasulo With Reorder buffer:



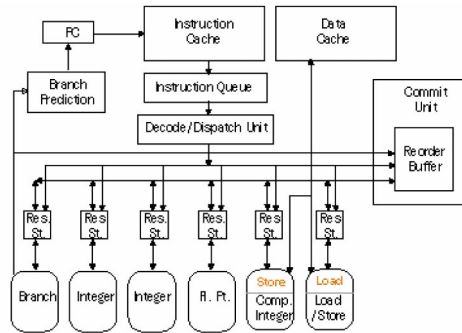
3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.52

## Dynamic Scheduling in PowerPC 604 and Pentium Pro

- Both In-order Issue, Out-of-order execution, In-order Commit



PPro central reservation station for any functional units with one bus shared by a branch and an integer unit

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.53

## Dynamic Scheduling in PowerPC 604 and Pentium Pro

| Parameter                          | PPC               | PPro |
|------------------------------------|-------------------|------|
| Max. instructions issued/clock     | 4                 | 3    |
| Max. instr. complete exec./clock   | 6                 | 5    |
| Max. instr. committed/clock        | 6                 | 3    |
| Instructions in reorder buffer     | 16                | 40   |
| Number of rename buffers           | 12 Int/8 FP       | 40   |
| Number of reservations stations    | 12                | 20   |
| No. integer functional units (FUs) | 2                 | 2    |
| No. floating point FUs             | 1                 | 1    |
| No. branch FUs                     | 1                 | 1    |
| No. complex integer FUs            | 1                 | 0    |
| No. memory FUs                     | 1 1 load +1 store |      |

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.54

## Dynamic Scheduling in Pentium Pro

- PPro doesn't pipeline 80x86 instructions
- PPro decode unit translates the Intel instructions into 72-bit micro-operations (- MIPS)
- Sends micro-operations to reorder buffer & reservation stations
- Takes 1 clock cycle to determine length of 80x86 instructions + 2 more to create the micro-operations
- Most instructions translate to 1 to 4 micro-operations
- Complex 80x86 instructions are executed by a conventional microprogram (8K x 72 bits) that issues long sequences of micro-operations

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.55

## Limits to Multi-Issue Machines

- Inherent limitations of ILP
  - 1 branch in 5: How to keep a 5-way superscalar busy?
  - Latencies of units: many operations must be scheduled
  - Need about Pipeline Depth x No. Functional Units of independent instructions to keep fully busy
  - Increase ports to Register File
    - VLIW example needs 7 read and 3 write for Int. Reg. & 5 read and 3 write for FP reg
  - Increase ports to memory
  - Current state of the art: Many hardware structures (such as issue/rename logic) has delay proportional to square of number of instructions issued/cycle

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.56

## Limits to ILP

- **Conflicting studies of amount**
  - Benchmarks (vectorized Fortran FP vs. integer C programs)
  - Hardware sophistication
  - Compiler sophistication
- **How much ILP is available using existing mechanisms with increasing HW budgets?**
- **Do we need to invent new HW/SW mechanisms to keep on processor performance curve?**
  - Intel MMX
  - Motorola AltaVec
  - Supersparc Multimedia ops, etc.

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.57

## Limits to ILP

Initial HW Model here; MIPS compilers.

Assumptions for ideal/perfect machine to start:

1. **Register renaming**—infinite virtual registers and all WAW & WAR hazards are avoided
2. **Branch prediction**—perfect; no mispredictions
3. **Jump prediction**—all jumps perfectly predicted => machine with perfect speculation & an unbounded buffer of instructions available
4. **Memory-address alias analysis**—addresses are known & a store can be moved before a load provided addresses not equal

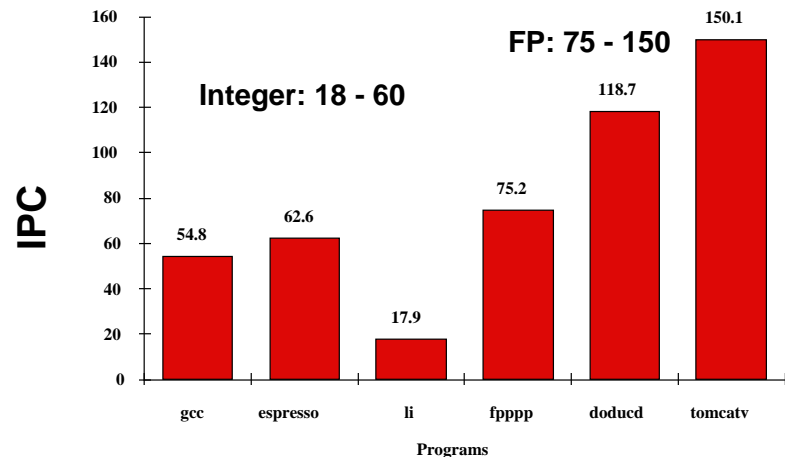
1 cycle latency for all instructions; unlimited number of instructions issued per clock cycle

3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.58

## Upper Limit to ILP: Ideal Machine

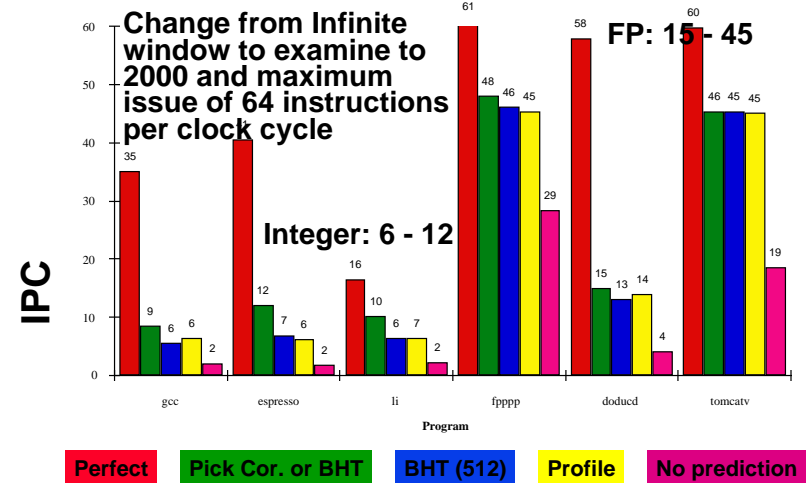


3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.59

## More Realistic HW: Branch Impact

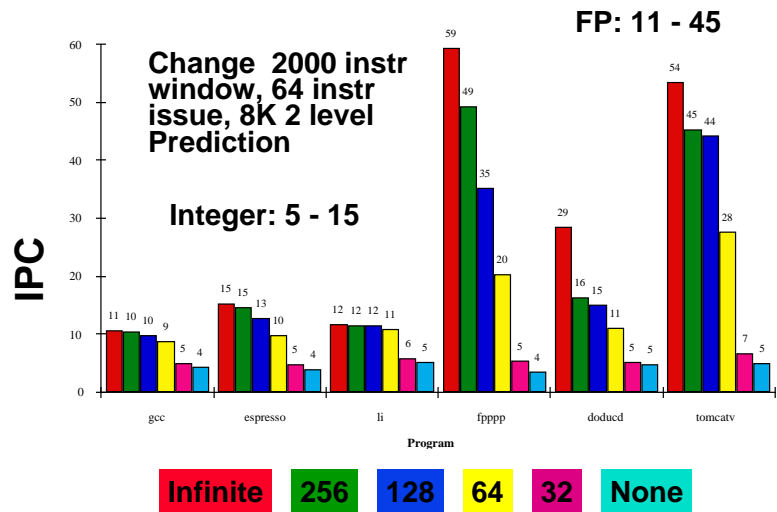


3/31/99

©UCB Spring 1999

CS152 / Kubiawicz  
Lec16.60

### More Realistic HW: Register Impact (rename regs)

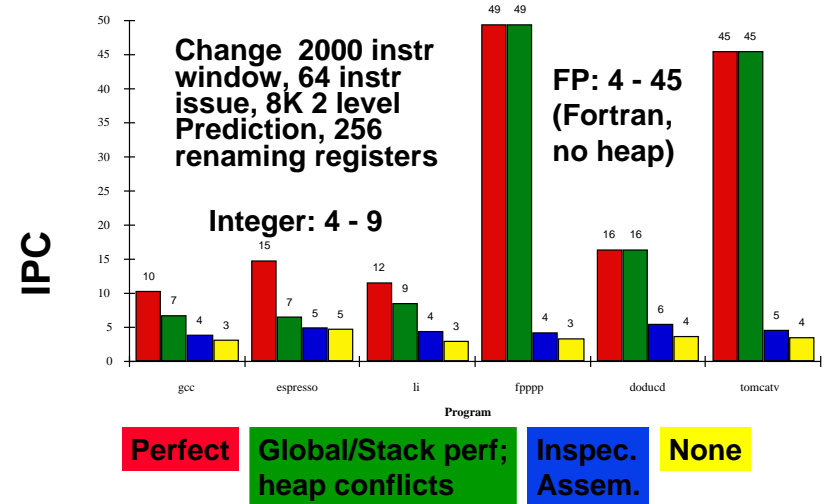


3/31/99

©UCB Spring 1999

CS152 / Kubiatiowicz  
Lec16.61

### More Realistic HW: Alias Impact

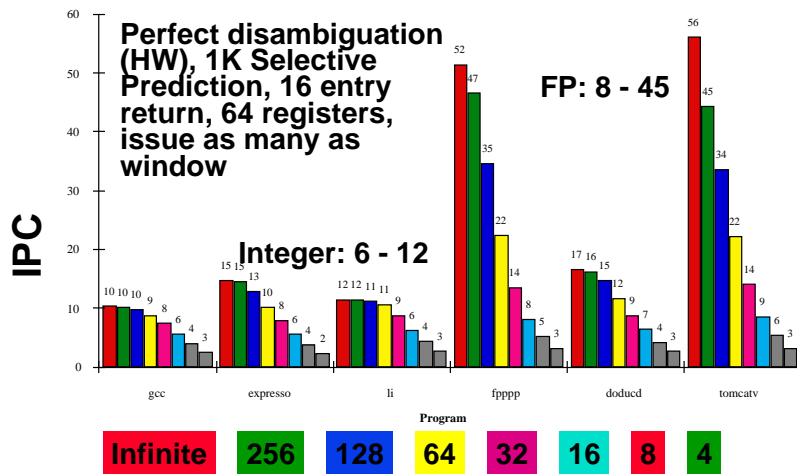


3/31/99

©UCB Spring 1999

CS152 / Kubiatiowicz  
Lec16.62

### Realistic HW for '9X: Window Impact



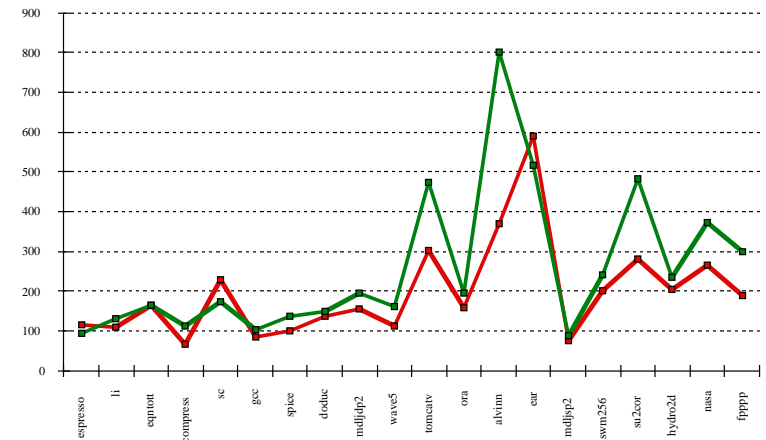
3/31/99

©UCB Spring 1999

CS152 / Kubiatiowicz  
Lec16.63

### Branic vs. Speed Demon(1993)

° 8-scalar IBM Power-2 @ 71.5 MHz (5 stage pipe)  
vs. 2-scalar Alpha @ 200 MHz (7 stage pipe)



3/31/99

Benchmark  
©UCB Spring 1999

CS152 / Kubiatiowicz  
Lec16.64

## Summary

- **Dynamic hardware schemes can unroll loops dynamically in hardware**
- **Branch prediction very important to good performance**
- **Precise exceptions/Speculation: Out-of-order execution, In-order commit (reorder buffer)**
- **Superscalar and VLIW:  $CPI < 1$  ( $IPC > 1$ )**
  - **Dynamic issue vs. Static issue**
  - **More instructions issue at same time  $\Rightarrow$  larger hazard penalty**
  - **Limitation is often number of instructions that you can successfully fetch and decode per cycle  $\Rightarrow$  "Flynn barrier"**
- **SW Pipelining**
  - **Symbolic Loop Unrolling to get most from pipeline with little code expansion, little overhead**