

Lecture 6: Instruction Level Parallelism 2: Loop-Level parallelism extraction, Data Flow, Explicit Register Renaming.

Prof. John Kubiawicz
Computer Science 252
Fall 1998

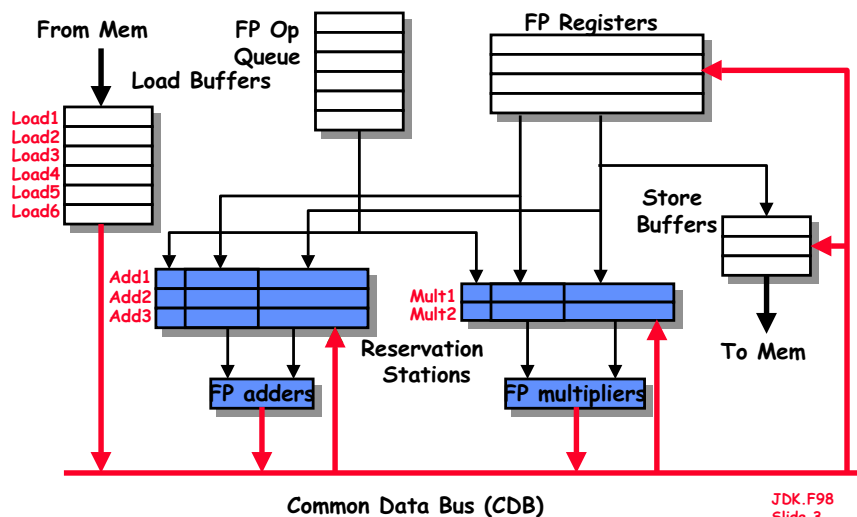
JDK.F98
Slide 1

Review: Dynamic hardware techniques 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

JDK.F98
Slide 2

Tomasulo Organization



JDK.F98
Slide 3

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

JDK.F98
Slide 4

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

JDK.F98
Slide 5

Loop Example

Instruction status:

ITER	Instruction	j	k	Issue	Comp	Result	Exec Write	Busy	Addr	Fu
1	LD	F0	0	R1				Load1	No	
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
0	80	Fu								

JDK.F98
Slide 6

Loop Example Cycle 1

Instruction status:

ITER	Instruction	j	k	Issue	Comp	Result	Exec Write	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							

JDK.F98
Slide 7

Loop Example Cycle 2

Instruction status:

ITER	Instruction	j	k	Issue	Comp	Result	Exec Write	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				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(F4)	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						

JDK.F98
Slide 8

Loop Example Cycle 3

Instruction status:

				Exec Write					
ITER	Instruction	j	k	Issue	CompResult	Busy	Addr	Fu	
1	LD	F0	0	R1	1	Yes	80		
1	MULTD	F4	F0	F2	2	No			
1	SD	F4	0	R1	3	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(F4)	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						

• Implicit renaming sets up "DataFlow" graph

JDK.F98
Slide 9

Loop Example Cycle 4

Instruction status:

				Exec Write					
ITER	Instruction	j	k	Issue	CompResult	Busy	Addr	Fu	
1	LD	F0	0	R1	1	Yes	80		
1	MULTD	F4	F0	F2	2	No			
1	SD	F4	0	R1	3	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(F4)	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

JDK.F98
Slide 10

Loop Example Cycle 5

Instruction status:

				Exec Write					
ITER	Instruction	j	k	Issue	CompResult	Busy	Addr	Fu	
1	LD	F0	0	R1	1	Yes	80		
1	MULTD	F4	F0	F2	2	No			
1	SD	F4	0	R1	3	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(F4)	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

JDK.F98
Slide 11

Loop Example Cycle 6

Instruction status:

				Exec Write					
ITER	Instruction	j	k	Issue	CompResult	Busy	Addr	Fu	
1	LD	F0	0	R1	1	Yes	80		
1	MULTD	F4	F0	F2	2	Yes	72		
1	SD	F4	0	R1	3	No			
2	LD	F0	0	R1	6	Yes	80	Mult1	
2	MULTD	F4	F0	F2	Store1	No			
2	SD	F4	0	R1	Store2	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(F4)	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

JDK.F98
Slide 12

Loop Example Cycle 7

Instruction status:

ITER	Instruction	j	k	Exec Write			Issue	CompResult	Busy	Addr	Fu
				S1	S2	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	Mult1	
2	MULTD	F4	F0	F2	7		Store2	No			
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
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
7	72	Fu	Load2	Mult2						

- Register file completely detached from computation
- First and Second iteration completely overlapped

JDK.F98
Slide 13

Loop Example Cycle 8

Instruction status:

ITER	Instruction	j	k	Exec Write			Issue	CompResult	Busy	Addr	Fu
				S1	S2	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	Mult1	
2	MULTD	F4	F0	F2	7		Store2	Yes	72	Mult2	
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
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						

JDK.F98
Slide 14

Loop Example Cycle 9

Instruction status:

ITER	Instruction	j	k	Exec Write			Issue	CompResult	Busy	Addr	Fu
				S1	S2	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	Mult1	
2	MULTD	F4	F0	F2	7		Store2	Yes	72	Mult2	
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
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

JDK.F98
Slide 15

Loop Example Cycle 10

Instruction status:

ITER	Instruction	j	k	Exec Write			Issue	CompResult	Busy	Addr	Fu
				S1	S2	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:
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

JDK.F98
Slide 16

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

JDK.F98
Slide 17

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?

JDK.F98
Slide 18

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						

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?

JDK.F98
Slide 20

JDK.F98
Slide 19

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?

JDK.F98
Slide 21

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						

JDK.F98
Slide 22

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						

JDK.F98
Slide 23

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						

JDK.F98
Slide 24

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	Code:	S1	S2	RS
Add1	No							LD	F0	0	R1
Add2	No							MULTD	F4	F0	F2
Add3	No							SD	F4	0	R1
Mult1	Yes	Multd						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						

JDK.F98
Slide 25

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	Code:	S1	S2	RS
Add1	No							LD	F0	0	R1
Add2	No							MULTD	F4	F0	F2
Add3	No							SD	F4	0	R1
Mult1	Yes	Multd						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						

JDK.F98
Slide 26

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
- Other idea: Tomasulo building "DataFlow" graph on the fly.

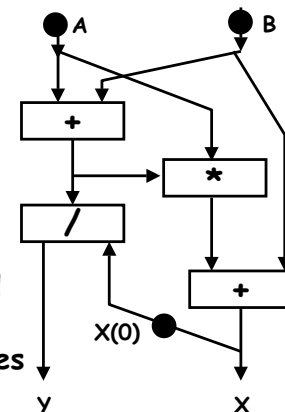
JDK.F98
Slide 27

Data-Flow Architectures

- Basic Idea: Hardware represents direct encoding of compiler dataflow graphs:

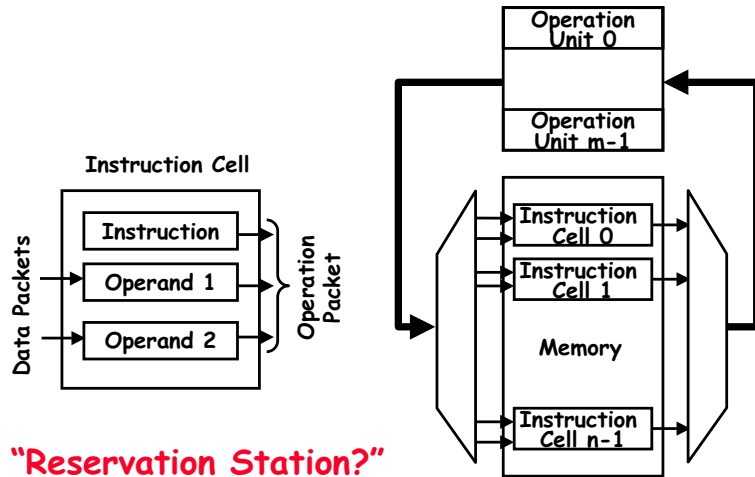
Input: a,b
 $y := (a+b)/x$
 $x := (a*(a+b))+b$
 output: y,x

- Data flows along arcs in "Tokens".
- When two tokens arrive at compute box, box "fires" and produces new token.
- Split operations produce copies of tokens



JDK.F98
Slide 28

Paper by Dennis and Misunas



"Reservation Station?"

JDK.F98
Slide 29

Brief, In-class discussion of Monsoon

JDK.F98
Slide 30

What about Precise Interrupts?

- Both Scoreboard and Tomasulo have:
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.

JDK.F98
Slide 31

Relationship between precise interrupts and speculation:

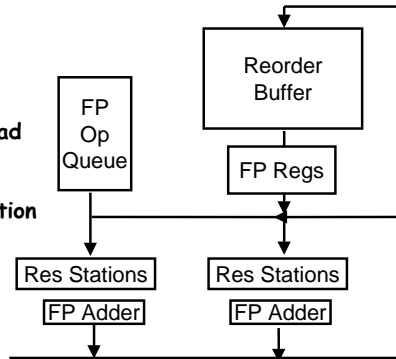
- Speculation is a form of guessing.
- Important for branch prediction:
 - Need to "take our best shot" at predicting branch direction.
 - If we issue multiple instructions per cycle, lose lots of potential instructions otherwise:
 - » Consider 4 instructions per cycle
 - » If take single cycle to decide 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*

JDK.F98
Slide 32

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**



JDK.F98
Slide 33

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)

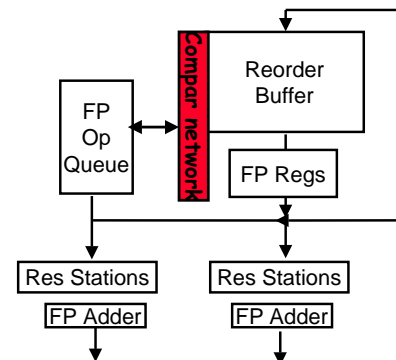
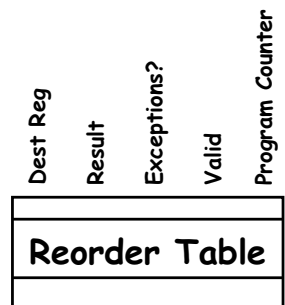
Write on Common Data Bus to all awaiting FUs & reorder buffer; mark reservation station available.

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 flushes reorder buffer (sometimes called "graduation")

JDK.F98
Slide 34

What are the hardware complexities with reorder buffer?



- Need fully-associative comparison with all reorder entries on instruction dispatch to find latest version of register.
- Need as many ports as register file

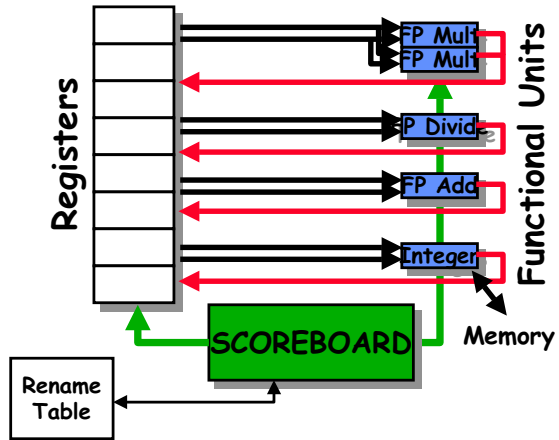
JDK.F98
Slide 35

Explicit Register Renaming

- Make use of a *physical* register file that is larger than number of registers specified by ISA
- Keep a translation table:
 - ISA register => physical register mapping
 - When register is written, replace table entry with new register from freelist.
 - Physical register becomes free when not being used by any instructions in progress.
- Pipeline can be exactly like "standard" DLX pipeline
 - IF, ID, EX, etc...
- Advantages:
 - Removes all WAR and WAW hazards
 - Like Tomasulo, good for allowing full out-of-order completion
 - Allows data to be fetched from a single register file
 - Makes speculative execution/precise interrupts easier:
 - » All that needs to be "undone" for precise break point is to undo the table mappings

JDK.F98
Slide 36

Question: Can we use explicit register renaming with scoreboard?



JDK.F98
Slide 37

Scoreboard Example

Instruction status:

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

Functional unit status:

Time Name	Busy	Op	dest		S1	S2	FU	FU	Fj?	Fk?
			Fi	Fj	Fk	Qj	Qk	Rj	Rk	
Int1	No									
Int2	No									
Multi	No									
Add	No									
Divide	No									

Register Rename and Result

Clock	FU									
	F0	F2	F4	F6	F8	F10	F12	...	F30	
1	P0	P2	P4	P6	P8	P10	P12		P30	

- Initialized Rename Table

JDK.F98
Slide 38

Renamed Scoreboard 1

Instruction status:

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

Functional unit status:

Time Name	Busy	Op	dest		S1	S2	FU	FU	Fj?	Fk?
			Fi	Fj	Fk	Qj	Qk	Rj	Rk	
Int1	Yes	Load	P32			R2				Yes
Int2	No									
Multi	No									
Add	No									
Divide	No									

Register Rename and Result

Clock	FU									
	F0	F2	F4	F6	F8	F10	F12	...	F30	
1	P0	P2	P4	P32	P8	P10	P12		P30	

- Each instruction allocates free register
- Similar to single-assignment compiler transformation

JDK.F98
Slide 39

Renamed Scoreboard 2

Instruction status:

Instruction	j	k	Read Exec Write		
			Issue	Oper	Comp Result
LD	F6	34+	R2	1	2
LD	F2	45+	R3	2	
MULTD	F0	F2	F4		
SUBD	F8	F6	F2		
DIVD	F10	F0	F6		
ADDD	F6	F8	F2		

Functional unit status:

Time Name	Busy	Op	dest		S1	S2	FU	FU	Fj?	Fk?
			Fi	Fj	Fk	Qj	Qk	Rj	Rk	
Int1	Yes	Load	P32			R2				Yes
Int2	Yes	Load	P34			R3				Yes
Multi	No									
Add	No									
Divide	No									

Register Rename and Result

Clock	FU									
	F0	F2	F4	F6	F8	F10	F12	...	F30	
2	P0	P34	P4	P32	P8	P10	P12		P30	

JDK.F98
Slide 40

Renamed Scoreboard 3

Instruction status:

Instruction	j	k	Read Exec Write		
			Issue	Oper	Comp Result
LD	F6	34+ R2	1	2	3
LD	F2	45+ R3	2	3	
MULTD	F0	F2 F4	3		
SUBD	F8	F6 F2			
DIVD	F10	F0 F6			
ADDD	F6	F8 F2			

Functional unit status:

Time Name	Busy	Op	dest		S1	S2	FU	FU	Fj?	Fk?
			Fi	Fj						
Int1	Yes	Load	P32			R2				Yes
Int2	Yes	Load	P34			R3				Yes
Multi1	Yes	Multd	P36	P34	P4	Int2			No	Yes
Add	No									
Divide	No									

Register Rename and Result

Clock		F0	F2	F4	F6	F8	F10	F12	...	F30
3	FU	P36	P34	P4	P32	P8	P10	P12		P30

JDK.F98
Slide 41

Renamed Scoreboard 4

Instruction status:

Instruction	j	k	Read Exec Write		
			Issue	Oper	Comp Result
LD	F6	34+ R2	1	2	3 4
LD	F2	45+ R3	2	3	4
MULTD	F0	F2 F4	3		
SUBD	F8	F6 F2	4		
DIVD	F10	F0 F6			
ADDD	F6	F8 F2			

Functional unit status:

Time Name	Busy	Op	dest		S1	S2	FU	FU	Fj?	Fk?
			Fi	Fj						
Int1	No									
Int2	Yes	Load	P34			R3				Yes
Multi1	Yes	Multd	P36	P34	P4	Int2			No	Yes
Add	Yes	Sub	P38	P32	P34			Int2	Yes	No
Divide	No									

Register Rename and Result

Clock		F0	F2	F4	F6	F8	F10	F12	...	F30
4	FU	P36	P34	P4	P32	P38	P10	P12		P30

JDK.F98
Slide 42

Renamed Scoreboard 5

Instruction status:

Instruction	j	k	Read Exec Write		
			Issue	Oper	Comp Result
LD	F6	34+ R2	1	2	3 4
LD	F2	45+ R3	2	3	4 5
MULTD	F0	F2 F4	3		
SUBD	F8	F6 F2	4		
DIVD	F10	F0 F6	5		
ADDD	F6	F8 F2			

Functional unit status:

Time Name	Busy	Op	dest		S1	S2	FU	FU	Fj?	Fk?
			Fi	Fj						
Int1	No									
Int2	No									
Multi1	Yes	Multd	P36	P34	P4				Yes	Yes
Add	Yes	Sub	P38	P32	P34				Yes	Yes
Divide	Yes	Divd	P40	P36	P32	Multi1			No	Yes

Register Rename and Result

Clock		F0	F2	F4	F6	F8	F10	F12	...	F30
5	FU	P36	P34	P4	P32	P38	P40	P12		P30

JDK.F98
Slide 43

Renamed Scoreboard 6

Instruction status:

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

Functional unit status:

Time Name	Busy	Op	dest		S1	S2	FU	FU	Fj?	Fk?
			Fi	Fj						
Int1	No									
Int2	No									
10 Multi1	Yes	Multd	P36	P34	P4				Yes	Yes
2 Add	Yes	Sub	P38	P32	P34				Yes	Yes
Divide	Yes	Divd	P40	P36	P32	Multi1			No	Yes

Register Rename and Result

Clock		F0	F2	F4	F6	F8	F10	F12	...	F30
6	FU	P36	P34	P4	P32	P38	P40	P12		P30

JDK.F98
Slide 44

Renamed Scoreboard 7

Instruction status:

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

Functional unit status:

Time Name	Busy	Op	dest		S1	S2	FU	FU	Fj?	Fk?
			Fi	Fj						
Int1	No									
Int2	No									
9 Mult1	Yes	Multd	P36	P34	P4				Yes	Yes
1 Add	Yes	Sub	P38	P32	P34				Yes	Yes
Divide	Yes	Divd	P40	P36	P32	Mult1		No	Yes	

Register Rename and Result

Clock		F0	F2	F4	F6	F8	F10	F12	...	F30
7	FU	P36	P34	P4	P32	P38	P40	P12		P30

JDK.F98
Slide 45

Renamed Scoreboard 8

Instruction status:

Instruction	j	k	Read		Exec		Write	
			Issue	Oper	Comp	Result		
LD	F6	34+	R2	1	2	3	4	
LD	F2	45+	R3	2	3	4	5	
MULTD	F0	F2	F4	3	6			
SUBD	F8	F6	F2	4	6		8	
DIVD	F10	F0	F6	5				
ADDD	F6	F8	F2					

Functional unit status:

Time Name	Busy	Op	dest		S1	S2	FU	FU	Fj?	Fk?
			Fi	Fj						
Int1	No									
Int2	No									
8 Mult1	Yes	Multd	P36	P34	P4				Yes	Yes
0 Add	Yes	Sub	P38	P32	P34				Yes	Yes
Divide	Yes	Divd	P40	P36	P32	Mult1		No	Yes	

Register Rename and Result

Clock		F0	F2	F4	F6	F8	F10	F12	...	F30
8	FU	P36	P34	P4	P32	P38	P40	P12		P30

JDK.F98
Slide 46

Renamed Scoreboard 9

Instruction status:

Instruction	j	k	Read		Exec		Write	
			Issue	Oper	Comp	Result		
LD	F6	34+	R2	1	2	3	4	
LD	F2	45+	R3	2	3	4	5	
MULTD	F0	F2	F4	3	6			
SUBD	F8	F6	F2	4	6	8	9	
DIVD	F10	F0	F6	5				
ADDD	F6	F8	F2					

Functional unit status:

Time Name	Busy	Op	dest		S1	S2	FU	FU	Fj?	Fk?
			Fi	Fj						
Int1	No									
Int2	No									
7 Mult1	Yes	Multd	P36	P34	P4				Yes	Yes
Add	No									
Divide	Yes	Divd	P40	P36	P32	Mult1		No	Yes	

Register Rename and Result

Clock		F0	F2	F4	F6	F8	F10	F12	...	F30
9	FU	P36	P34	P4	P32	P38	P40	P12		P30

JDK.F98
Slide 47

Renamed Scoreboard 10

Instruction status:

Instruction	j	k	Read		Exec		Write	
			Issue	Oper	Comp	Result		
LD	F6	34+	R2	1	2	3	4	
LD	F2	45+	R3	2	3	4	5	
MULTD	F0	F2	F4	3	6			
SUBD	F8	F6	F2	4	6	8	9	
DIVD	F10	F0	F6	5				
ADDD	F6	F8	F2	10				

Functional unit status:

Time Name	Busy	Op	dest		S1	S2	FU	FU	Fj?	Fk?
			Fi	Fj						
Int1	No									
Int2	No									
6 Mult1	Yes	Multd	P36	P34	P4				Yes	Yes
Add	Yes	Addd	P42	P38	P4				Yes	Yes
Divide	Yes	Divd	P40	P36	P32	Mult1		No	Yes	

Register Rename and Result

Clock		F0	F2	F4	F6	F8	F10	F12	...	F30
10	FU	P36	P34	P4	P42	P38	P40	P12		P30

- Notice that P32 not listed in Rename Table
- Still live. Must not be reallocated by accident

JDK.F98
Slide 48

Renamed Scoreboard 11

Instruction status:

Instruction	j	k	Read		Exec		Write	
			Issue	Oper	Comp	Result		
LD	F6	34+	R2	1	2	3	4	
LD	F2	45+	R3	2	3	4	5	
MULTD	F0	F2	F4	3	6			
SUBD	F8	F6	F2	4	6	8	9	
DIVD	F10	F0	F6	5				
ADDD	F6	F8	F2	10	11			

Functional unit status:

Time Name	Busy	Op	dest			FU	FU	Fj?	Fk?
			Fi	Fj	Fk				
Int1	No								
Int2	No								
5 Mult1	Yes	Multd	P36	P34	P4			Yes	Yes
2 Add	Yes	Addd	P42	P38	P34			Yes	Yes
Divide	Yes	Divd	P40	P36	P32	Mult1		No	Yes

Register Rename and Result

Clock		F0	F2	F4	F6	F8	F10	F12	...	F30
11	FU	P36	P34	P4	P42	P38	P40	P12		P30

JDK.F98
Slide 49

Renamed Scoreboard 12

Instruction status:

Instruction	j	k	Read		Exec		Write	
			Issue	Oper	Comp	Result		
LD	F6	34+	R2	1	2	3	4	
LD	F2	45+	R3	2	3	4	5	
MULTD	F0	F2	F4	3	6			
SUBD	F8	F6	F2	4	6	8	9	
DIVD	F10	F0	F6	5				
ADDD	F6	F8	F2	10	11			

Functional unit status:

Time Name	Busy	Op	dest			FU	FU	Fj?	Fk?
			Fi	Fj	Fk				
Int1	No								
Int2	No								
4 Mult1	Yes	Multd	P36	P34	P4			Yes	Yes
1 Add	Yes	Addd	P42	P38	P34			Yes	Yes
Divide	Yes	Divd	P40	P36	P32	Mult1		No	Yes

Register Rename and Result

Clock		F0	F2	F4	F6	F8	F10	F12	...	F30
12	FU	P36	P34	P4	P42	P38	P40	P12		P30

JDK.F98
Slide 49

Renamed Scoreboard 13

Instruction status:

Instruction	j	k	Read		Exec		Write	
			Issue	Oper	Comp	Result		
LD	F6	34+	R2	1	2	3	4	
LD	F2	45+	R3	2	3	4	5	
MULTD	F0	F2	F4	3	6			
SUBD	F8	F6	F2	4	6	8	9	
DIVD	F10	F0	F6	5				
ADDD	F6	F8	F2	10	11	13		

Functional unit status:

Time Name	Busy	Op	dest			FU	FU	Fj?	Fk?
			Fi	Fj	Fk				
Int1	No								
Int2	No								
3 Mult1	Yes	Multd	P36	P34	P4			Yes	Yes
0 Add	Yes	Addd	P42	P38	P34			Yes	Yes
Divide	Yes	Divd	P40	P36	P32	Mult1		No	Yes

Register Rename and Result

Clock		F0	F2	F4	F6	F8	F10	F12	...	F30
13	FU	P36	P34	P4	P42	P38	P40	P12		P30

JDK.F98
Slide 51

Renamed Scoreboard 14

Instruction status:

Instruction	j	k	Read		Exec		Write	
			Issue	Oper	Comp	Result		
LD	F6	34+	R2	1	2	3	4	
LD	F2	45+	R3	2	3	4	5	
MULTD	F0	F2	F4	3	6			
SUBD	F8	F6	F2	4	6	8	9	
DIVD	F10	F0	F6	5				
ADDD	F6	F8	F2	10	11	13	14	

Functional unit status:

Time Name	Busy	Op	dest			FU	FU	Fj?	Fk?
			Fi	Fj	Fk				
Int1	No								
Int2	No								
2 Mult1	Yes	Multd	P36	P34	P4			Yes	Yes
Add	No								
Divide	Yes	Divd	P40	P36	P32	Mult1		No	Yes

Register Rename and Result

Clock		F0	F2	F4	F6	F8	F10	F12	...	F30
14	FU	P36	P34	P4	P42	P38	P40	P12		P30

JDK.F98
Slide 52

Renamed Scoreboard 15

Instruction status:

Instruction	j	k	Read		Exec		Write	
			Issue	Oper	Comp	Result		
LD	F6	34+ R2	1	2	3	4		
LD	F2	45+ R3	2	3	4	5		
MULTD	F0	F2 F4	3	6				
SUBD	F8	F6 F2	4	6	8	9		
DIVD	F10	F0 F6	5					
ADDD	F6	F8 F2	10	11	13	14		

Functional unit status:

Time Name	Busy	Op	dest		S1	S2	FU	FU	Fj?	Fk?
			Fi	Fj						
Int1	No									
Int2	No									
1 Mult1	Yes	Multd	P36	P34	P4				Yes	Yes
Add	No									
Divide	Yes	Divd	P40	P36	P32	Mult1			No	Yes

Register Rename and Result

Clock		F0	F2	F4	F6	F8	F10	F12	...	F30
15	FU	P36	P34	P4	P42	P38	P40	P12		P30

JDK.F98
Slide 53

Renamed Scoreboard 16

Instruction status:

Instruction	j	k	Read		Exec		Write	
			Issue	Oper	Comp	Result		
LD	F6	34+ R2	1	2	3	4		
LD	F2	45+ R3	2	3	4	5		
MULTD	F0	F2 F4	3	6	16			
SUBD	F8	F6 F2	4	6	8	9		
DIVD	F10	F0 F6	5					
ADDD	F6	F8 F2	10	11	13	14		

Functional unit status:

Time Name	Busy	Op	dest		S1	S2	FU	FU	Fj?	Fk?
			Fi	Fj						
Int1	No									
Int2	No									
0 Mult1	Yes	Multd	P36	P34	P4				Yes	Yes
Add	No									
Divide	Yes	Divd	P40	P36	P32	Mult1			No	Yes

Register Rename and Result

Clock		F0	F2	F4	F6	F8	F10	F12	...	F30
16	FU	P36	P34	P4	P42	P38	P40	P12		P30

JDK.F98
Slide 54

Renamed Scoreboard 17

Instruction status:

Instruction	j	k	Read		Exec		Write	
			Issue	Oper	Comp	Result		
LD	F6	34+ R2	1	2	3	4		
LD	F2	45+ R3	2	3	4	5		
MULTD	F0	F2 F4	3	6	16	17		
SUBD	F8	F6 F2	4	6	8	9		
DIVD	F10	F0 F6	5					
ADDD	F6	F8 F2	10	11	13	14		

Functional unit status:

Time Name	Busy	Op	dest		S1	S2	FU	FU	Fj?	Fk?
			Fi	Fj						
Int1	No									
Int2	No									
Mult1	No									
Add	No									
Divide	Yes	Divd	P40	P36	P32	Mult1			Yes	Yes

Register Rename and Result

Clock		F0	F2	F4	F6	F8	F10	F12	...	F30
17	FU	P36	P34	P4	P42	P38	P40	P12		P30

JDK.F98
Slide 55

Renamed Scoreboard 18

Instruction status:

Instruction	j	k	Read		Exec		Write	
			Issue	Oper	Comp	Result		
LD	F6	34+ R2	1	2	3	4		
LD	F2	45+ R3	2	3	4	5		
MULTD	F0	F2 F4	3	6	16	17		
SUBD	F8	F6 F2	4	6	8	9		
DIVD	F10	F0 F6	5	18				
ADDD	F6	F8 F2	10	11	13	14		

Functional unit status:

Time Name	Busy	Op	dest		S1	S2	FU	FU	Fj?	Fk?
			Fi	Fj						
Int1	No									
Int2	No									
Mult1	No									
Add	No									
40 Divide	Yes	Divd	P40	P36	P32	Mult1			Yes	Yes

Register Rename and Result

Clock		F0	F2	F4	F6	F8	F10	F12	...	F30
18	FU	P36	P34	P4	P42	P38	P40	P12		P30

JDK.F98
Slide 56

Explicit Renaming Support Includes:

- Rapid access to a table of translations
- A physical register file that has more registers than specified by the ISA
- Ability to figure out which physical registers are free.
 - No free registers \Rightarrow stall on issue
- Thus, register renaming doesn't require reservation stations. However:
 - Many modern architectures use explicit register renaming + Tomasulo-like reservation stations to control execution.

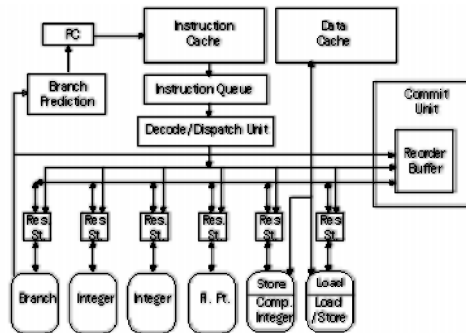
JDK.F98
Slide 57

Discussion of R10000

JDK.F98
Slide 58

Dynamic Scheduling in PowerPC 604 and Pentium Pro

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



Pentium Pro more like a scoreboard since central control vs. distributed

JDK.F98
Slide 59

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
Window (Instrs in reorder buffer)	16	40
Number of reservations stations	12	20
Number of rename registers	8int/12FP	40
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	

Q: How pipeline 1 to 17 byte x86 instructions?

JDK.F98
Slide 60

Dynamic Scheduling in Pentium Pro

- PPro doesn't pipeline 80x86 instructions
- PPro decode unit translates the Intel instructions into 72-bit micro-operations (- DLX)
- 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
- 12-14 clocks in total pipeline (- 3 state machines)
- Many 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

JDK.F98
Slide 61

Getting CPI < 1: Issuing Multiple Instructions/Cycle

- Two variations
- **Superscalar**: varying no. instructions/cycle (1 to 8), scheduled by compiler or by HW (Tomasulo)
 - IBM PowerPC, Sun UltraSparc, DEC Alpha, HP 8000
- **(Very) Long Instruction Words (V)LIW**: fixed number of instructions (4-16) scheduled by the compiler; put ops into wide templates
 - Joint HP/Intel agreement in 1999/2000?
 - Intel Architecture-64 (IA-64) 64-bit address
 - Style: "Explicitly Parallel Instruction Computer (EPIC)"
- Anticipated success lead to use of **Instructions Per Clock cycle (IPC)** vs. CPI

JDK.F98
Slide 62

Getting CPI < 1: Issuing Multiple Instructions/Cycle

- Superscalar DLX: 2 instructions, 1 FP & 1 anything else
 - Fetch 64-bits/clock cycle; Int on left, FP on right
 - Can only issue 2nd instruction if 1st instruction issues
 - More ports for FP registers to do FP load & FP op in a pair

Type	Pipe Stages						
Int. instruction	IF	ID	EX	MEM	WB		
FP instruction	IF	ID	EX	MEM	WB		
Int. instruction		IF	ID	EX	MEM	WB	
FP instruction		IF	ID	EX	MEM	WB	
Int. instruction			IF	ID	EX	MEM	WB
FP instruction			IF	ID	EX	MEM	WB

- 1 cycle load delay expands to **3 instructions** in SS
 - instruction in right half can't use it, nor instructions in next slot

JDK.F98
Slide 63

Review: Unrolled Loop that Minimizes Stalls for Scalar

```

1 Loop: LD      F0,0(R1)           LD to ADDD: 1 Cycle
2          LD      F6,-8(R1)       ADDD to SD: 2 Cycles
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

JDK.F98
Slide 64

Loop Unrolling in Superscalar

	Integer instruction	FP instruction	Clock cycle
Loop:	LD F0,0(R1)		1
	LD F6,-8(R1)		2
	LD F10,-16(R1)	ADDD F4,F0,F2	3
	LD F14,-24(R1)	ADDD F8,F6,F2	4
	LD F18,-32(R1)	ADDD F12,F10,F2	5
	SD 0(R1),F4	ADDD F16,F14,F2	6
	SD -8(R1),F8	ADDD F20,F18,F2	7
	SD -16(R1),F12		8
	SD -24(R1),F16		9
	SUBI R1,R1,#40		10
	BNEZ R1,LOOP		11
	SD -32(R1),F20		12

- Unrolled 5 times to avoid delays (+1 due to SS)
- 12 clocks, or 2.4 clocks per iteration (1.5X)

JDK.F98
Slide 65

Multiple Issue Challenges

- While Integer/FP split is simple for the HW, get CPI of 0.5 only for programs with:
 - Exactly 50% FP operations
 - No hazards
- If more instructions issue at same time, greater difficulty of decode and issue
 - Even 2-scalar => examine 2 opcodes, 6 register specifiers, & decide if 1 or 2 instructions can issue
- VLIW: 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

JDK.F98
Slide 66

Loop Unrolling in VLIW

Memory reference 1	Memory reference 2	FP operation 1	FP op. 2	Int. op/branch	Clock
LD F0,0(R1)	LD F6,-8(R1)				1
LD F10,-16(R1)	LD F14,-24(R1)				2
LD F18,-32(R1)	LD F22,-40(R1)	ADDD F4,F0,F2	ADDD F8,F6,F2		3
LD F26,-48(R1)		ADDD F12,F10,F2	ADDD F16,F14,F2		4
		ADDD F20,F18,F2	ADDD F24,F22,F2		5
SD 0(R1),F4	SD -8(R1),F8	ADDD F28,F26,F2			6
SD -16(R1),F12	SD -24(R1),F16				7
SD -32(R1),F20	SD -40(R1),F24			SUBI R1,R1,#48	8
SD -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)

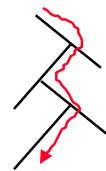
Average: 2.5 ops per clock, 50% efficiency

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

JDK.F98
Slide 67

Trace Scheduling

- Parallelism across IF branches vs. LOOP branches
- Two steps:
 - Trace Selection
 - » Find likely sequence of basic blocks (*trace*) of (statically predicted or profile predicted) long sequence of straight-line code
 - Trace Compaction
 - » Squeeze trace into few VLIW instructions
 - » Need bookkeeping code in case prediction is wrong
- Compiler undoes bad guess (discards values in registers)
- Subtle compiler bugs mean wrong answer vs. poorer performance; no hardware interlocks



JDK.F98
Slide 68

Advantages of HW (Tomasulo) vs. SW (VLIW) Speculation

- HW determines address conflicts
- HW better branch prediction
- HW maintains precise exception model
- HW does not execute bookkeeping instructions
- Works across multiple implementations
- SW speculation is much easier for HW design

JDK.F98
Slide 69

Superscalar v. VLIW

- Smaller code size
- Binary compatibility across generations of hardware
- Simplified Hardware for decoding, issuing instructions
- No Interlock Hardware (compiler checks?)
- More registers, but simplified Hardware for Register Ports (multiple independent register files?)

JDK.F98
Slide 70

Intel/HP "Explicitly Parallel Instruction Computer (EPIC)"

- 3 Instructions in 128 bit "groups"; field determines if instructions dependent or independent
 - Smaller code size than old VLIW, larger than x86/RISC
 - Groups can be linked to show independence > 3 instr
- 64 integer registers + 64 floating point registers
 - Not separate files per functional unit as in old VLIW
- Hardware checks dependencies (interlocks => binary compatibility over time)
- Predicated execution (select 1 out of 64 1-bit flags)
=> 40% fewer mispredictions?
- IA-64 : name of instruction set architecture; EPIC is type
- Merced is name of first implementation (1999/2000?)
- LIW = EPIC?

JDK.F98
Slide 71

Dynamic Scheduling in Superscalar

- Dependencies stop instruction issue
- Code compiler for old version will run poorly on newest version
 - May want code to vary based on version of architecture

JDK.F98
Slide 72

Dynamic Scheduling in Superscalar

- How to issue two instructions and keep in-order instruction issue for Tomasulo?
 - Assume 1 integer + 1 floating point
 - 1 Tomasulo control for integer, 1 for floating point
- Issue 2X Clock Rate, so that issue remains in order
- Only FP loads might cause dependency between integer and FP issue:
 - Replace load reservation station with a load queue; operands must be read in the order they are fetched
 - Load checks addresses in Store Queue to avoid RAW violation
 - Store checks addresses in Load Queue to avoid WAR, WAW
 - Called "decoupled architecture"

JDK.F98
Slide 73

Performance of Dynamic SS

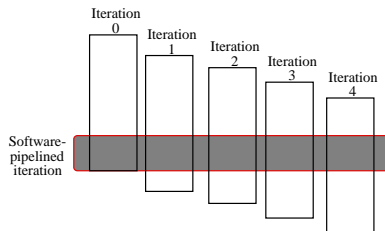
Iteration no.	Instructions	Issues	Executes	Writes result
			clock-cycle number	
1	LD F0,0(R1)	1	2	4
1	ADDD F4,F0,F2	1	5	8
1	SD 0(R1),F4	2	9	
1	SUBI R1,R1,#8	3	4	5
1	BNEZ R1,LOOP	4	5	
2	LD F0,0(R1)	5	6	8
2	ADDD F4,F0,F2	5	9	12
2	SD 0(R1),F4	6	13	
2	SUBI R1,R1,#8	7	8	9
2	BNEZ R1,LOOP	8	9	

- 4 clocks per iteration; only 1 FP instr/iteration
- Branches, Decrements issues still take 1 clock cycle
- How get more performance?

JDK.F98
Slide 74

Software Pipelining

- Observation: if iterations from loops are independent, then can get more ILP by taking instructions from different iterations
- Software pipelining: reorganizes loops so that each iteration is made from instructions chosen from different iterations of the original loop (- Tomasulo in SW)



JDK.F98
Slide 75

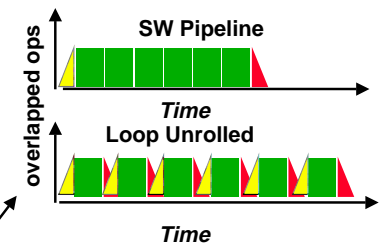
Software Pipelining Example

Before: Unrolled 3 times After: Software Pipelined

1 LD F0,0(R1)	1 SD 0(R1),F4 ; Stores M[i]
2 ADDD F4,F0,F2	2 ADDD F4,F0,F2 ; Adds to M[i-1]
3 SD 0(R1),F4	3 LD F0,-16(R1); Loads M[i-2]
4 LD F6,-8(R1)	4 SUBI R1,R1,#8
5 ADDD F8,F6,F2	5 BNEZ R1,LOOP
6 SD -8(R1),F8	
7 LD F10,-16(R1)	
8 ADDD F12,F10,F2	
9 SD -16(R1),F12	
10 SUBI R1,R1,#24	
11 BNEZ R1,LOOP	

• Symbolic Loop Unrolling

- Maximize result-use distance
- Less code space than unrolling
- Fill & drain pipe only once per loop vs. once per each unrolled iteration in loop unrolling



JDK.F98
Slide 76

Limits to Multi-Issue Machines

- Inherent limitations of ILP
 - 1 branch in 5: How to keep a 5-way VLIW busy?
 - Latencies of units: many operations must be scheduled
 - Need about Pipeline Depth x No. Functional Units of independent Difficulties in building HW
 - Easy: More instruction bandwidth
 - Easy: Duplicate FUs to get parallel execution
 - Hard: Increase ports to Register File (bandwidth)
 - » VLIW example needs 7 read and 3 write for Int. Reg. & 5 read and 3 write for FP reg
 - Harder: Increase ports to memory (bandwidth)
 - Decoding Superscalar and impact on clock rate, pipeline depth?

JDK.F98
Slide 77

Limits to Multi-Issue Machines

- Limitations specific to either Superscalar or VLIW implementation
 - Decode issue in Superscalar: how wide practical?
 - VLIW code size: unroll loops + wasted fields in VLIW
 - » IA-64 compresses dependent instructions, but still larger
 - VLIW lock step => 1 hazard & all instructions stall
 - » IA-64 not lock step? Dynamic pipeline?
 - VLIW & binary compatibility IA-64 promises binary compatibility

JDK.F98
Slide 78

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?

JDK.F98
Slide 79

Limits to ILP

Initial HW Model here; MIPS compilers.

Assumptions for ideal/perfect machine to start:

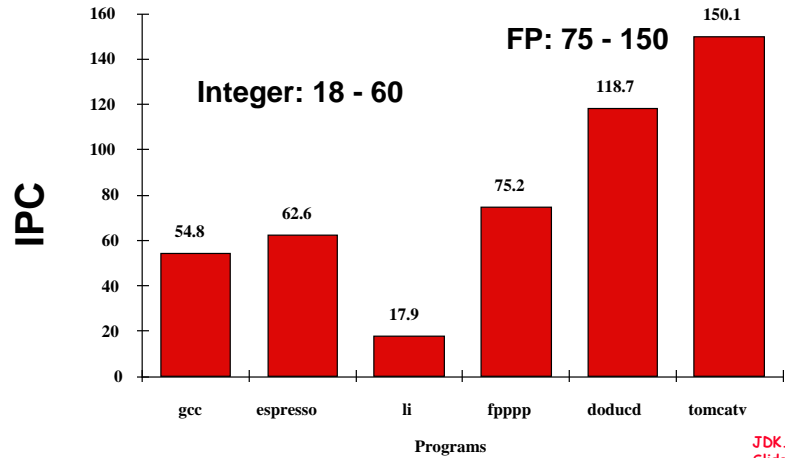
1. **Register renaming**-infinite virtual registers and all WAW & WAR hazards are avoided
2. **Branch prediction**-perfect; no mispredictions
3. **Jump prediction**-all jumps perfectly predicted => machine with perfect speculation & an unbounded buffer of instructions available
4. **Memory-address alias 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

JDK.F98
Slide 80

Upper Limit to ILP: Ideal Machine

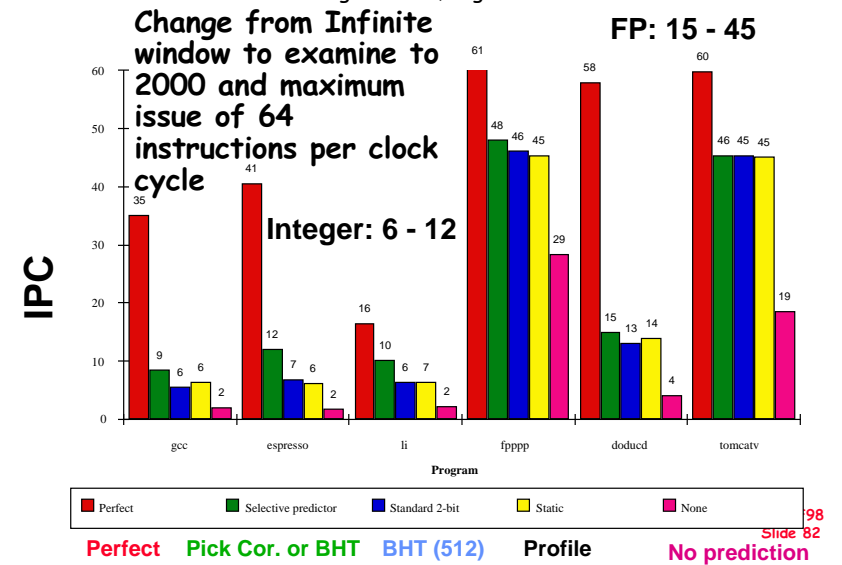
(Figure 4.38, page 319)



JDK.F98
Slide 81

More Realistic HW: Branch Impact

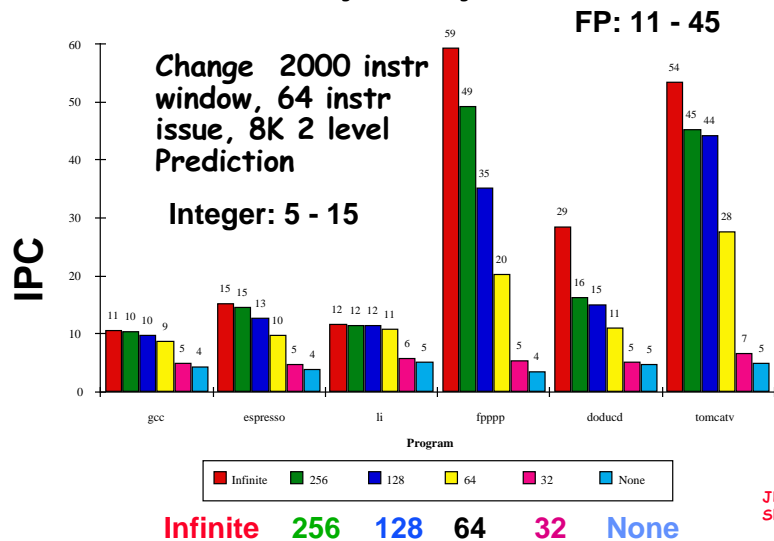
Figure 4.40, Page 323



JDK.F98
Slide 82

More Realistic HW: Register Impact

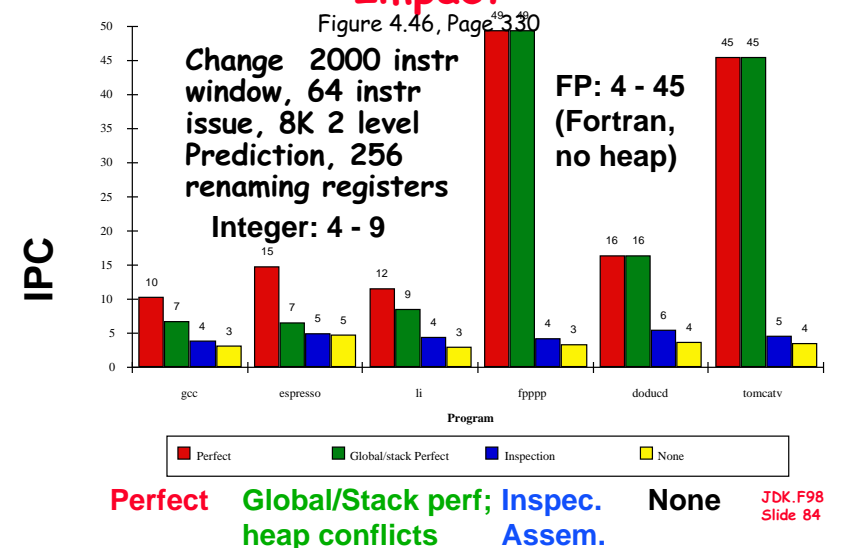
Figure 4.44, Page 328



JDK.F98
Slide 83

More Realistic HW: Alias Impact

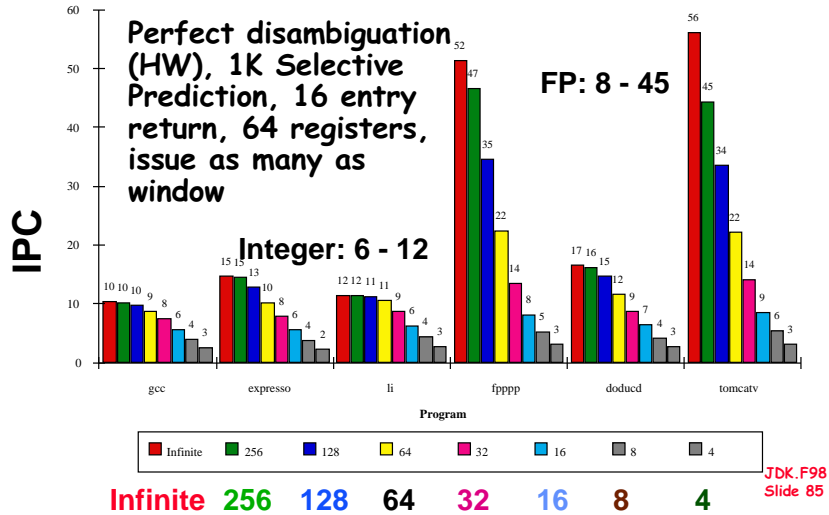
Figure 4.46, Page 330



JDK.F98
Slide 84

Realistic HW for '9X: Window Impact

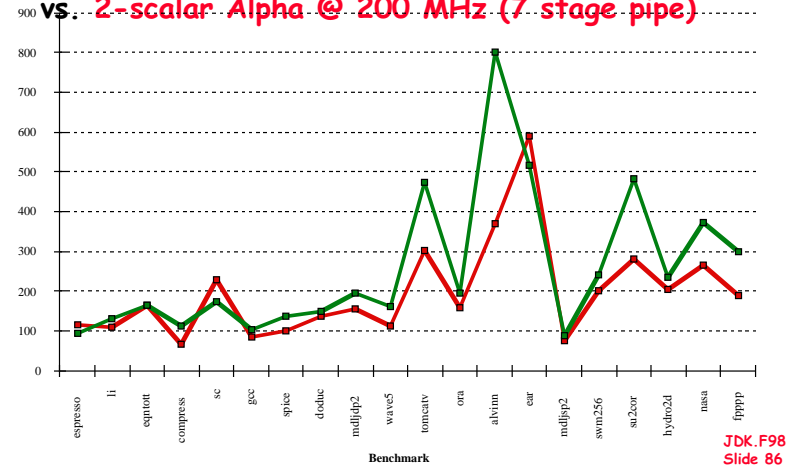
(Figure 4.48, Page 332)



Braniac vs. Speed Demon(1993)

• 8-scalar IBM Power-2 @ 71.5 MHz (5 stage pipe)

vs. 2-scalar Alpha @ 200 MHz (7 stage pipe)

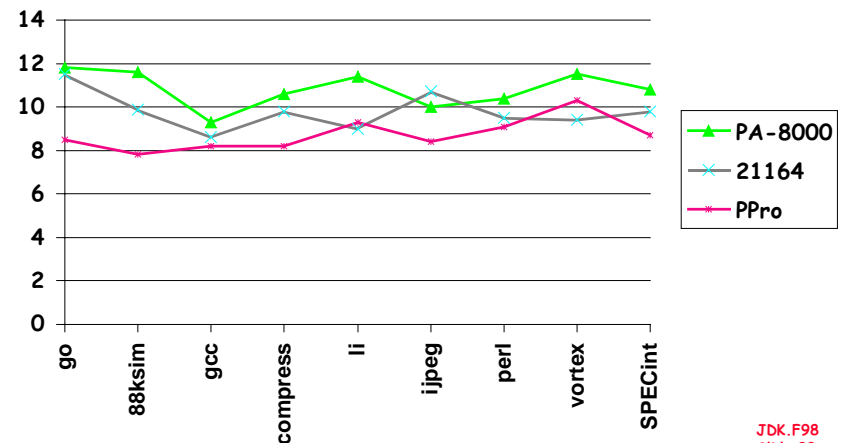


3 1996 Era Machines

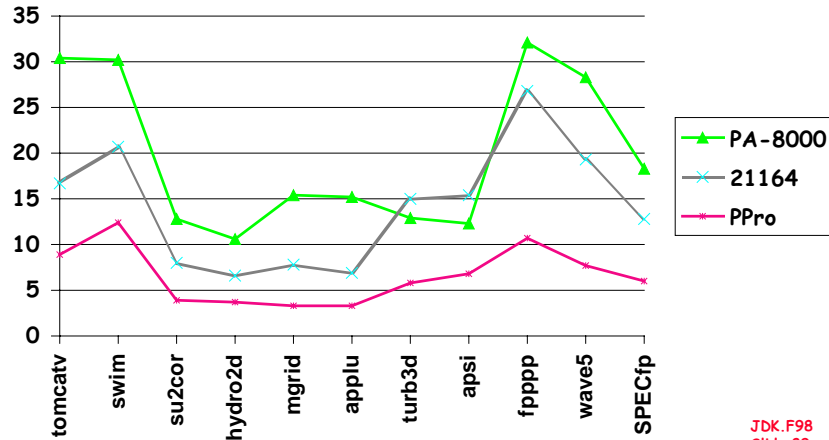
	Alpha 21164	PPro	HP PA-8000
Year	1995	1995	1996
Clock	400 MHz	200 MHz	180 MHz
Cache	8K/8K/96K/2M	8K/8K/0.5M	0/0/2M
Issue rate	2int+2FP	3 instr (x86)	4 instr
Pipe stages	7-9	12-14	7-9
Out-of-Order	6 loads	40 instr (μ op)	56 instr
Rename regs	none	40	56

JDK.F98 Slide 87

SPECint95base Performance (July 1996)



SPECfp95base Performance (July 1996)



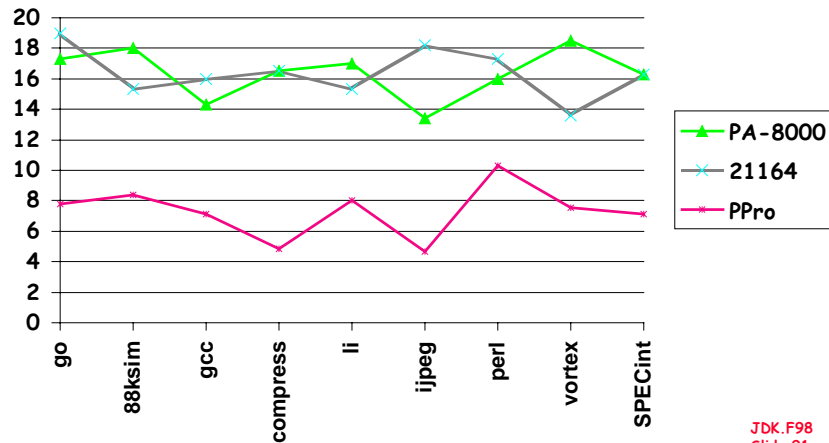
JDK.F98
Slide 89

3 1997 Era Machines

	Alpha 21164	Pentium II	HP PA-8000
Year	1995	1996	1996
Clock	<u>600 MHz ('97)</u>	<u>300 MHz ('97)</u>	<u>236 MHz</u>
Cache	8K/8K/96K/2M	<u>16K/16K/0.5M</u>	<u>0/0/4M</u>
Issue rate	2int+2FP	3 instr (x86)	4 instr
Pipe stages	7-9	12-14	7-9
Out-of-Order loads		40 instr (μ op)	56 instr
Rename regs	none	40	56

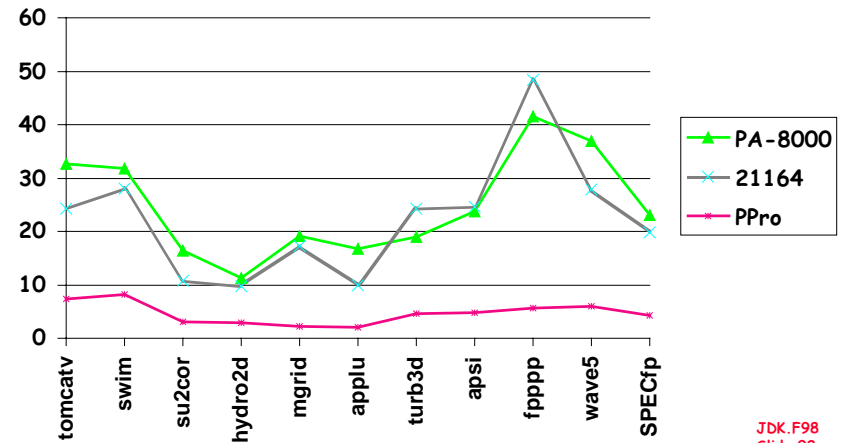
JDK.F98
Slide 90

SPECint95base Performance (Oct. 1997)



JDK.F98
Slide 91

SPECfp95base Performance (Oct. 1997)



JDK.F98
Slide 92

Summary

- Dynamic hardware schemes can unroll loops dynamically in hardware
- Explicit Renaming: more physical registers than needed by ISA. Uses a translation table
- 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 => larger hazard penalty
- SW Pipelining
 - Symbolic Loop Unrolling to get most from pipeline with little code expansion, little overhead