

Patent Number:

Date of Patent:

[11]

[45]

United States Patent [19]

Chamdani et al.

[54] DISTRIBUTED INSTRUCTION QUEUE

- [75] Inventors: Joseph I. Chamdani, Marietta; Cecil O. Alford, Lawrenceville, both of Ga.
- [73] Assignee: Georgia Tech Research Corp., Atlanta, Ga.
- [21] Appl. No.: 08/489,509
- [22] Filed: Jun. 12, 1995

[56] References Cited

U.S. PATENT DOCUMENTS

12/1975	Eaton et al 395/421.09
2/1988	Shonai et al 395/585
4/1988	Shintani et al 395/395
6/1988	Shonai et al 395/800.23
10/1988	Torii et al 395/800
1/1990	Yamaguchi et al 395/583
2/1991	Cocke et al 395/393
9/1991	McLogan et al 395/678
6/1992	Strehler 365/49
7/1992	Johnson 395/375
8/1992	Johnson 395/375
5/1993	Wilson et al 395/275
11/1993	Jouppi et al 395/425
9/1994	Tran 395/393
10/1994	Shebanow et al 395/394
11/1994	Levitan 395/800
12/1994	Iadonato et al 364/491
5/1995	Saito et al 395/375
4/1997	Ando 395/855
	12/1975 2/1988 4/1988 6/1988 10/1988 1/1990 2/1991 9/1991 6/1992 7/1992 8/1992 5/1993 11/1993 9/1994 10/1994 12/1994 5/1995 4/1997

OTHER PUBLICATIONS

Sohi, "Instruction Issue Logic for High–Performance, Interruptible, Multiple Function Unit, Pipelined Computers", 1990 IEEE, pp. 349–359. "An Efficient Algorithm for Exploring Multiple Arithmetic Units," Tomasulo IBM Journal, Jan. 1967, pp. 25–33.

6,112,019

Aug. 29, 2000

"Implementation of Precise Interrupts In Pipelined Processors," James E. Smith Andrew R. Pleszkun,© 1985 IEEE, pp. 36–44.

"Instruction Issue Logic in Pipelined Supercomputers", Shlomo Weiss & James E. Smith, © 1984 IEEE, Transactions on Computers, vol. c-33, No. 11, pp. 1012–1022 Nov. 1984.

"The Metaflow Architecture", Popescu et al, IEEE Micro, © 1991 IEEE, pp. 10–13, 63–73.

Primary Examiner-David Y. Eng

Attorney, Agent, or Firm—Thomas, Kayden, Horstemeyer & Risley, L.L.P.

[57] ABSTRACT

A distributed instruction queue (DIQ) in a superscalar microprocessor supports multi-instruction issue, decoupled data flow scheduling, out-of-order execution, register renaming, multi-level speculative execution, and precise interrupts. The DIQ provides distributed instruction shelving without storing register values, operand value copying, and result value forwarding, and supports in-order issue as well as out-of-order issue within its functional unit. The DIQ allows a reduction in the number of global wires and replacement with private-local wires in the processor. The DIQ's number of global wires remains the same as the number of DIQ entries and data size increases. The DIQ maintains maximum machine parallelism and the actual performance of the microprocessor using the DIQ is better due to reduced cycle time or more operations executed per cycle.

22 Claims, 40 Drawing Sheets









Register Mapping Table in IBM RS/6000 Floating-Point Unit



Note: N = the length of longest functional-unit pipeline

Fig. 3

Reorder Buffer Organization

	tag	opcode	S1	α(S1)	S2	α(S2)	D	α(D)	β(D)	12
top										
	I6	fadd	R0	2	R4	2	R0	2	2	8
	15	fadd	R4	1	R6	1	R4	1	1	4
	I4	fadd	R6	0	R 7	0	R6	0	0	0
	13	fadd	R4	0	_R5	0	R4	0	0	0
1	12	fadd	R 0	1	R2	1	R0	1	1	4
	I1_	fadd	R2	0	R 3	0	R2	0	0	0
oottom	10	fadd	R0	0	_R1	0	R0	0	0	0

(a) Before Issue

	tag	opcode	S1 (α(S1)	S2	α(S2)	D	α(D)	β(D)	²			
top				ometri			. for						
	subsequent instructions												
					ļ								
	16	fadd	R0	1	R4	1	R0	1	1	4			
	15	fadd	R4	0	R6	0	R4	0	0	0			
bottom	I2	fadd	R0	0	R2	0	R0	0	0	0			

(b) After Issue and Completion

- Note: S1/S2 = first/second source register identifier, D = destination register identifier,
 - $\alpha(X) = \#$ of times register X is designated as a destination register in preceding inst (below it),
 - $\beta(X) = \#$ of times register X is designated as a source register in preceding instruction (below it),

= issue index = $\alpha(S1) + \alpha(S2) + \alpha(D) + \beta(D)$.

8-Entry Dispatch Stack

ູດງ	/cle —	<u>► 1</u>	2	3 4	5	6	7	8	9_1	0,1	11_	12	13	14	15	11	.12	13	.14	15	
alloc	ate —	<u>▶</u>][4	15	<u>I6</u>	┨──╁			_				_			<u> </u>			
IS writeb	sue — ack —			<u>11 µ3</u>	114	+					<u>_</u>	5	10	\rightarrow		12	<u> </u>	15	16		
re	ack etire	┢	++		+	+				<u>3 </u> 1	4 2 1		-+	-+		12	12		10	7	
				L	-	<u> </u>				<u>, 11</u>	<u>S I</u>	41					112		10	0	1
to complete after issu														issu	cles ied.						
	Entry	PC	Opcode	Sourc	<u>æ 0</u>	peran	id 1	Sourc	ce Oj	pera	ind 2	De	stina	ation	Exe	ecu	ted l	Exce	p-		
	Number			ready	tag	conte	ent	ready	/ tag	cor	nten	¶tag	COL	ntent				tion	5		
tail>	7								1			 									
	6	I6	fadd	0	0.2	-		0	4.2		-	0.3	<u> </u>	•		0		•	_ ●	– allo)C
	5	I5	fadd	0	4.1	-		0	6.1		-	4.2		-		0		-	_ ◄	- noi	t rdy
	4	I4	fadd	1	-	R6	0	1	-	F	7 7	6.1		-		0		-	_ ◄	– 4th	issue
	3	13	fadd	1	-	R4	0	1	-	F	२5 ⁰	4.1		-		0		-	_ •	– 3rc	l issue
	2	I2	fadd	0	0.1	-		0	2.1		-	0.2	2	-		0		-	•	- noi	t rdy
	1	I1	fadd	1	-	R2	0	1	-	F	₹3 ⁰	2.1		-		0		-]•	– 2n	d issue
head →	0	10	fadd	1	-	R0	0	1	-	F	र1 ⁰	0.1		-		0		-	_⊦	- 1st	issue
		(b)) RUU	Sna	psh	not a	t C	ycle	7	N	lote	: F r	γi ^k egi	mea ster	ans r R	s th i (r	ne k egi:	(th i ster	nst rei	anc nam	e of ling)
	Entry	PC	Opcode	Sourc	e O	peran	id 1	Sourc	e Op	bera	nd 2	De	stina	ation	Exe	ecul	ted	Exce	p-1		
	Number			ready	tag	conte	ent	ready	tag	cor	nten	tag	cor	ntent				tions	5		
tail	7																				
	6	16	fadd	0	0.2	-		0	4.2		-	0.3		-		0		•	_ ◄	– noi	: rdy
	5	15	fadd	0	4.1	-		0	6.1		-	4.2		-		0		-	_ ∙	- not	t rdy
	4	14	fadd	1	-	R6	0	1	-	F	77 ⁰	6.1		-		0		-	•	- in e	exec
	3	13	fadd	1	-	R4	0	1	-	F	ر 59	4.1		-		0		-	•	— in (exec
	2	12	fadd	1	0.1	R0	1	1	2.1	F	۲2 ¹	0.2	2	-		0		•		– rea	idy
	1	I1	fadd	1	-	R2	0	1	-	F	२ 30	2.1	F	R21		1		non	• •	– wrl	back
head 🗕	0	10	fadd	1		R0	0	1	-	F	۲1 ⁰	0.1	F	r0 ¹		1		non	•	– ret	ire
		(0	:) RUl	J Sna	aps	hot a	at C	Cycle	9												

Register Update Unit



Metaflow Architecture

(cycle —	▶	1	2	3	4	1	5	6	7		8	9	10		11
allo	oćate –	► 10, 11,	, I2, I3 I	4,15	, I6											
i	ssue –	→		_I 0 , I	1 13,	I 4			I2	I	5		_			16
write	back –	◆						IO, I1	I3, I4	1			I2	I5		
1	retire –	◆							I0, I1	I3,	I 4			12		13
								(a) Ins	truction	Timing						
	N	ote: Assi	ume the	re are	e 4 alloca	te ports,	4 ret	ire ports	, 							
1		2 10	ating-po	ant a	aa FUS (<u>class nun</u>	12) M	lith 3-cy	cle laten	су.						-
	Index	Source	Operan	d 1	Source	Operand	12	De	estinatio	n	Disp.	FU	Exec.	Oncode	PC	
	Index	Locked	R num	ID	Locked	R num	ID	latest	Rnum	content		class		opoode		
►	7															
	6	1	0	0.2	1	4	0.5	1	0	_	0	2	0	fadd	16] ←_ alloc
	5	1	4	0.3	1	6	0.4	1	4	_	0	2	0	fadd	15	l → alloc
	4	0	6	_	0	7	_	1	6		0	2	0	fadd	I4	📥 alloc
	3	0	4		0	5	_	0	4	_	0	2	0	fadd	13	📥 ready
	2	1	0	0.0	1	2	0.1	0	0	_	0	2	0	fadd	I2	I ← not rdy
	1	0	2	_	0	3	_	1	2	_	1	2	0	fadd	I 1	- 1st iss
h	0	0	0	_	0	1	_	0	0	_	1	2	0	fadd	10	¯ ⊸ 1st iss
							(b)	DRIS S	napshot	t at Cvcl	e 2	Note:	Ri ^k m	eans the	k th ins	tance of
							(~)						registe	er Ri (reg	ister re	enaming)
		Source	Operan	d 1	Source	Operan	d 2	De	estinatio	n	Diam	FU	-		D C	7
	Index	Locked	Locked R num ID		Locked	Rnum	ID	latest	R num	content	Disp.	class	Exec.	Opcode	PC	
i•	7						[_			-
	6	1	0	0.2	1	4	0.5	1	0	_	0	2	0	fadd	16	- not rdy
	5	0	4	0.3	0	6	0.4	1	4	_	0	2	0	fadd	15	- ready
	4	0	6		0	7	1_	1	6	R61	1	2	1	fadd	I4	∣ ← wrtbk
	3	0	4		Ō	5		Ó	4	R ₄ 1	1	2	1	fadd	13	wrtbk
h	2	0	0	0.0	Ō	2	0.1	Ō	i o		1	2	Ó	fadd	12	- I → issue
	1	0	2		0	3		1	2	 	1	2	1	fadd	11	- retire
	0	0	0		0	1	<u> </u>	0	0		1	2		fadd	10	- retire
				L		L	<u> </u>		inanehol	t at Cycle	<u> </u>	·	<u> </u>			
		Fio	1.7			Deferred	رت) محمد	adulina	Register	-Ronam	una Inc	truction	n Shalf	۱		
		· · · · · · · · · · · · · · · · · · ·			DRIS (Deletted	-scne	eauiing,	Register	-Renam	ing ins	ructio	1 Shell)		





IBM RS/6000 Architecture

Fig. 9

	Cycle — 1 2		2 3	3 4	5	6	7	8	9	10	11	12	13	14	15	
Branch	ſIF	I 7 ⁰ I I 8 ⁰ I I 9 ⁰ I I10 ⁰ I	11 ⁰ I1 12 ⁰ 13 ⁰ 14 ⁰	15 ⁰	17 ¹ 18 ¹ 19 ¹ 110 ¹	111 ¹ 112 ¹ 113 ¹ 114 ¹	I15 ¹									
Unit	Disp/BRE		170 I 180 I	9 ⁰ I11 ⁰ 10 ⁰ I12 ⁰ (I15 ⁰) 1130 0 1140 1) 1150	17 ¹ 18 ¹	19 ¹ 110 ¹	I11 ¹ I12 ¹ (I15 ¹)	I13 ¹ I14 ¹ I15 ¹	:						
Fixed- Point -	Fixed- Point FXE			7° 19° 8° 110° 17°	I11 ⁰ I12 ⁰ I9 ⁰ I7 ⁰	I13 ⁰ I14 ⁰ I12 ⁰ I9 ⁰	171 181 114 ⁰	19 110 17 ¹	I11 ¹ I12 ¹ I9 ¹ I7 ¹	I13 ¹ I14 ¹ I12 ¹ I9 ¹	I14 ¹	 				
	(PD		I	7 ⁰ I9 ⁰ 8 ⁰ I10 ⁰ 17 ⁰	111 ⁰ 112 ⁰ 19 ⁰	I13 ⁰ I14 ⁰ I11 ⁰	171 181 113 ⁰	19 ¹ 110 ¹ 171	I11 ¹ I12 ¹ I9 ¹	I13 ¹ I14 ¹ I11 ¹	I13 ¹					
Floating- Point Unit	Remap FPD FPE1 FPE2 EDWP			<u>18</u> 0	110 ⁰ 18 ⁰	112 ⁰ 110 ⁰ 18 ⁰	I14 ⁰ I11 ⁰ I10 ⁰ I8 ⁰	I8 ¹ I13 ⁰ I11 ⁰ I10 ⁰ I81	110 ¹ 181 113 ⁰ 111 ⁰	1121 1101 181 1130 1110	I141 I111 I101 I81 I130	I13 ¹ I11 ¹ I10 ¹ I8 ¹	I131 I111 I101	<u>13</u> 1	131	
Pipe	line Stage	Note	IF Disp	= b/BRE =	Fetch 4 Dispate	instru h fixea	uctions d-and ter ins	from t floating tructio	the l-c g-poin ns.	ache a t instru	irrays, ictions	place to FX	them U and	into B I FPU.	U's I-b Exect	uffers. ute branch and
Cycle-by) = = = = nap =	 =Decode a fixed-point instruction and access the integer RF for operands. =Execute a fixed-point instruction. For load/store instruction, compute its address =Access the D-cache arrays. Data is returned to either the FXU or FPU. =Write back result value to the integer register file. Data for load is written in this cycle. =Pre-decode instructions in preparation for register renaming. =Map architectural registers of a floating-point instruction to physical registers. 														
in the	ibm RS/6 ig. 10	000	FPC FPE FPE FPV) = 1 = 2 = VB =	Decode Execut Secono Write b	e a floa e a flo l cycle ack re	ating-p ating-p of MA sult va	oint in ooint in AF pipe alue to	structi Istruct eline. the flo	on and ion. Th pating-	d acce nis is tl point F	ss the ne first RF.	floatir t cycle	ng-poil of MA	nt RF f AF pipe	or operands aline.











HP PA7100 Architecture



Sheet 16 of 40





Intel Pentium Architecture

Fig. 17

6,112,019

U.S. Patent

Aug. 29, 2000

Sheet 17 of 40



floating-point result buses

A Superscaler Processor w/ Reorder Buffer and Reservation Stations



Reorder Buffer (RB) Implemented with True FIFO Array



Result-Forwarding Hazard







Aug. 29, 2000

Sheet 23 of 40

6,112,019



tail(i)_sel[k]= 1 means that the dest reg num and tag from dispatch bus i is routed /allocated to MRB_Cell[k]

R_alloc_en[i]=1 means that there is an allocation to "R" MRB from inst dispatch bus i head(j)_sel[k]=1 means that MRB_Cell(k) is currently in the j^{th} retire port valid_ret[j]=means that the j^{th} retire port contains an MRB entry that is ready to retire

Support Logic for the MRB





Notes: i ∈[0,N_d-1], k ∈[0,N_{fag}-1]

tail(i)_sel[k] =1 means that dispatch bus i is routed/allocated to FAQ_Cell[k] head_sel[k]=1 means that FAQ_Cell[k] is currently the bottom entry FAQ_alloc_en[i]=1 means that there is an cond branch allocation to FAQ from inst dispatch bus i

Support Logic for the RBFAQ





Aug. 29, 2000

Sheet 28 of 40



head_sel[k] =1 means that DIQ_Cell[k] is currently the bottom instruction entry DIQ_alloc_en[i]=1 means that there is an allocation to this DIQ unit from inst dispatch base





Aug. 29, 2000

Sheet 31 of 40







Aug. 29, 2000

Sheet 34 of 40






Aug. 29, 2000

Sheet 37 of 40

6,112,019



Fig. 38

6,112,019

700 —	*				
	710	720	730	740	
	n_empty	last	reg_num	reg_tag]
n_empty last reg_num reg_tag	= set when entry = set when entry = register numbe = register tag, ins	is allocate is allocate r, correspo t_ID of the	d, reset whe d, reset whe onds to regist producer ins	n entry is r n there is a er file add strucion	etired a new assignment to register reg num ress/location

RTRB ENTRY

Fig. 39



DISTRIBUTED INSTRUCTION QUEUE

BACKGROUND OF THE INVENTION

1. Field of the Invention

This invention relates generally to an apparatus and a method for improving processor microarchitecture in superscalar microprocessors. In particular, the invention relates to an apparatus and a method for a modified reorder buffer and a distributed instruction queue that increases the efficiency by reducing the hardware complexity, execution time, and the number of global wires in superscalar microprocessors that support multi-instruction issue, decoupled dataflow scheduling, out-of-order execution, register renaming, multi-level speculative execution, load bypassing, and precise interrupts.

2. Background of the Related Art

The main driving force in the research and development of microprocessor architectures is improving performance/ unit cost. The true measure of performance is the time 20 (seconds) required to execute a program. The execution time of a program is basically determined by three factors (see Patterson and Hennessey, Computer Architecture: A Quantitative Approach, Morgan Kaufmann Publishers, 1990); the number of instructions executed in the program (dynamic 25 Inst_Count), the average number of clock cycles per instruction (CPI), and the processing cycle time (Clock_ Period), or

To improve performance (reduce execution time), it is necessary to reduce one or more factors. The obvious one to reduce is Clock Period, by means of semiconductor/VLSI technology improvements such as device scaling, faster circuit structures, better routing techniques, etc. A second 35 weak-memory ordering model. Relaxed ordering allows an approach to performance improvement is architecture design. CISC and VLIW architectures take the approach of reducing Inst_Count. RISC and superscalar architectures attempt to reduce the CPI. Superpipelined architectures increase the degree of pipelining to reduce the Clock 40 Period.

The true measure of cost is dollars/unit to implement and manufacture a microprocessor design in silicon. This hardware cost is driven by many factors such as die size, die yield, wafer cost, die testing cost, packaging cost, etc. The 45 architectural choices made in a microprocessor design affect all these factors.

It is desirable to focus on finding microarchitecture techniques/alternatives to improve the design of superscalar microprocessors. The term microprocessor refers to a pro- 50 regions of code concurrently (multiple flows of control). An cessor or CPU that is implemented in one or a small number of semiconductor chips. The term superscalar refers to a microprocessor implementation that increases performance by concurrent execution of scalar instructions, the type of instructions typically found in general-purpose micropro- 55 complexities and limitations involved in parallelizing a cessors. It should be understood that hereinafter, the term "processor" also means "microprocessor".

A superscalar architecture can be generalized as a processor architecture that fetches and decodes multiple scalar instructions from a sequential, single-flow instruction 60 stream, and executes them concurrently on different functional units. In general, there are seven basic processing steps in superscalar architectures; fetch, decode, dispatch, issue, execute, writeback, and retire. FIG. 1 illustrates these basic steps. 65

First, multiple scalar instructions are fetched simultaneously from an instruction cache/memory or other storage

unit. Current state-of-the-art superscalar microprocessors fetch two or four instructions simultaneously. Valid fetched instructions (the ones that are not after a branch-taken instruction) are decoded concurrently, and dispatched into a central instruction window (FIG. 1a) or distributed instruction queues or windows (FIG. 1b). Shelving of these instructions is necessary because some instructions cannot execute immediately, and must wait until their data dependencies and/or resource conflicts are resolved. After an instruction is 10 ready it is issued to the appropriate functional unit. Multiple ready instructions are issued simultaneously, achieving parallel execution within the processor. Execution results are written back to a result buffer first. Because instructions can complete out-of-order and speculatively, results must be 15 retired to register file(s) in the original, sequential program order. An instruction and its result can retire safely if it completes without an exception and there are no exceptions or unresolved conditional branches in the preceding instructions. Memory stores wait at a store buffer until they can commit safely.

The parallel executions in superscalar processors demands high memory bandwidth for instructions and data. Efficient instruction bandwidth can be achieved by aligning and merging the decode group. Branching causes wasted decoder slots on the left side (due to unaligned branch target addresses) and on the right side (due to a branch-taken instruction that is not at the end slot). Aligning shifts branch target instructions to the left most slot to utilize all decoder slots. Merging fills the slots to the right of a branch-taken instruction with the branch target instructions, combining different instruction runs into one dynamic instruction stream. Efficient data bandwidth can be achieved by load bypassing and load forwarding (M. Johnson, Superscalar Microprocessor Design, Prentice-Hall, 1991), a relaxed or out-of-order sequence of reads and writes, to optimize the use of the data bus. Stores to memory cannot commit until they are safe (retire step). Forcing loads and stores to commence in order will delay the loads significantly and stall other instructions that wait on the load data. Load bypassing allows a load to bypass stores in front of it (out-of-order execution), provided there is no read-afterwrite hazard. Load forwarding allows a load to be satisfied directly from the store buffer when there is a read-after-write dependency. Executing loads early is safe because load data is not written directly to the register file.

Classic superscalar architectures accomplish fine-grain parallel processing at the instruction level, which is limited to a single flow of control. They cannot execute independent instruction stream external to superscalar processors appears the same as in CISC or RISC uniprocessors; a sequential, single-flow instruction stream. It is internally that instructions are distributed to multiple processing units. There are sequential, single-flow instruction stream. The following six superscalar features-multi-instruction issue, decoupled dataflow scheduling, out-of-order execution, register renaming, speculative execution, and precise interrupts-are key in achieving this goal. They help improve performance and ensure correctness in superscalar processors.

Multi-instruction issue is made possible by widening a conventional, serial processing pipeline in the "horizontal" direction to have multiple pipeline streams. In this manner multiple instructions can be issued simultaneously per clock cycle. Thus, superscalar microprocessors must have multiple execution/functional units with independent pipeline

streams. Also, to be able to sustain multi-instruction issue at every cycle, superscalar microprocessors fetch and decode multiple instructions at a time.

Decoupled dataflow scheduling is supported by buffering all decoded instructions into an instruction window(s), before they are scheduled for execution. The instruction window(s) essentially "decouples" the decode and execute stage. There are two primary objectives. The first is to maintain the flow of instruction fetching and decoding by not forcing a schedule of the decoded instructions right away. This reduces unnecessary stalls. Instructions are allowed to take time to resolve data dependencies and/or resource conflicts. The second is to improve the look-ahead capability of the processor. With the instruction window, a processor is now able to look ahead beyond the stalled instructions to discover others that are ready to execute. The issue logic includes a dependency check to allow an instruction to "fire" or execute as soon as its operands are available and its resource conflicts are resolved. Unlike sequential Von Neumann machines, the control hardware does not have to sequence each instruction and decide explicitly when it can 20 execute. This is the essence of dataflow scheduling,

Out-of-order execution helps reduce instruction stalls due to data dependencies, bypassing the stalled or incomplete instructions. There are three types of out-of-order execution, categorized by their aggressiveness: (a) in-order issue with 25 out-of-order completion, (b) partial out-of-order issue with out-of-order completion, and (c) full out-of-order issue with out-of-order completion. The first type always issues instructions sequentially, in the original program order, but they can complete out-of-order due to different latencies or 30 stages in some functional units' pipelines. The second type restricts instruction issue to be in order only within a functional unit, but can be out of order amongst multiple functional units. The third type allows full out-of-order issue tional units.

Register renaming is necessary to eliminate the side effects of out-of-order execution, i.e., artificial dependencies on registers-those dependencies other than true date dependency (read-after-write hazard). There are two types of 40 challenges in superscalar architecture design are to find: an artificial dependencies, anti dependency (write-after-read hazard) and output dependency (write-after-write hazard) (M. Johnson, Superscalar Microprocessor Design, Prentice-Hall, 1991). They are caused by register-set limitations. The compiler's register allocation process minimizes the register 45 usage by reusing registers as much as possible. This action blurs the distinction between register and value. Register renaming effectively reintroduces the distinction by renaming the registers in hardware, creating a new instance of a register for each new register assignment.

Speculative execution avoids stalls and reduces the penalty due to control dependencies. For every conditional branch, a superscalar processor predicts the likely branch direction, with help from software (static branch prediction) or hardware (dynamic branch prediction). Instructions from 55 Improve Design of Superscalar Microprocessors," Copythe predicted path are fetched and executed speculatively, without waiting for the outcome of the branch test. By scheduling instructions across multiple, unresolved conditional branches (multi-level speculative execution), more instruction parallelism is potentially extracted, improving 60 the processor's performance. Due to the speculative nature, some conditional branches may be incorrectly predicted. A mechanism to recover and restart must be provided so that correct results can still be produced in the event of mispredicted branches. Recovery cancels the effect of instructions 65 processed under false predictions, and restart reestablishes the correct instruction sequence.

Precise interrupts are supported to guarantee the correct processor state before servicing the interrupt. Out-of-order execution complicates the restarting of an interrupted program. At the time an exception is detected, some instructions beyond the exception instruction might have been executed, as a result of allowing out-of-order execution. The effects on registers and memory by any instructions beyond the precise-repair point [?] must be nullified or repaired before going to the interrupt handler routine. The hardware support for precise interrupts should not be too costly if there is already hardware support for speculative execution.

There are two key microarchitecture elements in superscalar hardware that determine the success in achieving the above goal, result shelving and instruction shelving. Result 15 shelving is the key to support register renaming, out-oforder execution, speculative execution, and precise interrupts. Instruction shelving is the key to support multiinstruction issue, decoupled dataflow scheduling, and outof-order execution. Review of the literature suggests that the reorder buffer (RB) is the most complete result shelving technique (see, for example U.S. Pat. Nos. 5,136,697 to Johnson and No. 5,345,569 to Tran for discussions of conventional reorder buffers), and the reservation station (RS) is the best instruction shelving technique to give maximum machine parallelism. However, these two techniques have implementation drawbacks. The RB requires associative lookup that must be prioritized during each operand read. This results in relatively complex and slow circuit implementation. Also, the RB requires substantial shared-global buses for its operand and result buses, and the need to use dummy branch entries to support speculative execution which increases the RB entry usage. The RS requires tremendous amounts of shared (heavily-loaded), global (chip-wide) wires to support its operand value copywithin a functional unit as well as amongst multiple func- 35 ing and result value forwarding. With increasingly smaller transistor sizes, the dominant factors in determining silicon area and propagation delays is not the transistor, but metal wire, especially the ones that run across or all over the chip.

With the many promises that lie ahead, the research efficient utilization of the vast chip real-estate, the highspeed transistors, and the available instruction parallelism. The hardware improvements that lead to enhanced performance must be coupled with compiler/software scheduling improvements, however. There is a need for these improvements to be cost effective, or, at best, to actually reduce the cost of a superscalar microprocessor while increasing efficiency. In accordance with the above, we should avoid the tendency to design an overly complex superscalar architec-50 ture that produces mediocre gains which could have been easily achieved by compiler optimizations or that are cost limiting.

The present invention is discussed at length in the doctoral dissertation entitled "Microarchitecture Techniques to right© 1995, Georgia Institute of Technology, of one of the co-inventors, Joseph I. Chamdani, the subject matter of which is incorporated herein by reference. Hereinafter, the above dissertation will be referred to as Chamdani's dissertation.

This invention addresses architectural improvements to the design of superscalar processors that support the six key superscalar features. The primary objective of the invention was to find a better design alternative to the reservation station technique (considered the best known distributed instruction shelving technique to give maximum machine parallelism). The superscalar technique invented is: the

50

Distributed Instruction Queue (DIQ). The DIQ is a new distributed instruction shelving technique that offers a significantly more efficient (i.e., better performance/cost) implementation than the reservation station (RS) technique by eliminating operand value/copying and result value for- 5 warding.

The DIQ shelving technique offers a more efficient (i.e., good performance/cost) implementation of distributed instruction windows by eliminating the two major implementation drawbacks in the RS technique, operand value 10 copying and result forwarding. The DIQ can support in-order issue as well as out-of-order issue within its functional unit. The cost analysis suggests an improvement in almost every hardware component, with major reductions in Chamdani's dissertation). The expensive shared-global wires are mostly replaced by private-local wires that are easier to route, have less propagation delay, and occupy much smaller silicon area. The DIQ's number of global wires remains the same as the number of DIQ entries and 20 data size increase. A performance analysis using cycle-bycycle simulators confirms that the good characteristics of the RS technique in achieving maximum machine parallelism have been maintained in the DIQ technique (see Chamdani's dissertation). The out-of-order DIQ technique is at par with 25 the RS technique in terms of cycle-count performance, but higher in terms of overall performance if the improved clock frequency is factored in. The in-order issue DIQ sacrifices slightly on the cycle-count performance, which can easily be recovered through faster and simpler circuit implementation. 30 In the end, the actual speed or performance of a processor using the DIQ technique is faster due to reduced cycle time or more operations executed per cycle.

One object of the invention is to provide an improved superscalar processor.

Another object of the invention is to provide a distributed instruction queue that does not store register values.

A further object of the invention is to eliminate the need for operand value copying in a superscalar microprocessor.

Yet another object of the invention is to eliminate the need 40 for result value forwarding in a superscalar processor.

One other object of the invention is to provide a processor having reduced global buses.

One advantage of the invention is that it can improve the speed of a superscalar processor.

Another advantage of the invention is that it can reduce the amount of global buses required in a superscalar processor.

A further advantage of the invention is that it can allow for issuing of instructions in any order.

Still another advantage of the invention is that it can support multi-level speculative execution.

One feature of the invention is that it includes local bus architecture between register units and functional units.

These and other objects, advantages, and features are 55 provided by a distributed instruction queue, comprising: at least one entry cell having at least one entry field; at least one allocate port, each of the at least one allocate port connected to each of the at least one entry cell for allocation of a decoded instruction to the at least one entry cell; an issue 60 port connected to a predetermined one of the at least one entry cell, wherein instructions are issued through the issue port under logic control in any order from one of the at least one entry cell and the distributed instruction queue stores no register value.

Implementations of the invention may include one or more of the following features: a tail pointer logic unit to

determine the correct tail pointer position of each of the at least one allocate port; a head pointer logic unit to adjust a head pointer to point to a predetermined one of the at least one entry cell; an issue pointer logic unit to adjust an issue pointer to point to the one of the at least one entry cell for issuing of the instructions; the distributed instruction queue eliminates operand value copying and result value forwarding; the distributed instruction queue is operated independently of any other distributed instruction queue; the instructions are issued in-order; the instruction are issued out-oforder; and the instructions are issued in some form of limited out-of-order issue.

These above and other objects, advantages, and features of the invention will become more apparent from the folthe use of global wires, comparators, and multiplexers (see 15 lowing description thereof taken in conjunction with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 shows basic steps in superscalar architectures.

FIG. 2 shows a register mapping table in the IBM RS/6000 floating point unit.

FIG. 3 illustrates reorder buffer organization.

FIG. 4 shows an 8-entry dispatch stack.

FIG. 5 depicts a register update unit.

FIG. 6 is a diagram of the Metaflow architecture.

FIG. 7 shows the DRIS (Deferred-Scheduling, Register-Renaming Instruction Shelf).

FIG. 8 shows Tomasulo's Reservation Stations at IBM 360/91.

FIG. 9 is a diagram of the IBM RS/6000 architecture.

FIG. 10 shows cycle-by-cycle execution in the IBM RS/6000.

FIG. 11 is a diagram of the Lightning SPARC architec-35 ture.

FIG. 12 is a diagram of the SuperSPARC architecture.

FIG. 13 is a diagram of the DEC Alpha 21064 architecture.

FIG. 14 shows the DECchip 21064 pipeline stages.

FIG. 15 shows the HP PA-7100 architecture.

FIG. 16 shows the HP PA-7100 pipeline stages.

FIG. 17 shows the Intel Pentium architecture.

FIG. 18 is a diagram of a superscalar processor with reorder buffer and reservation stations.

FIG. 19 depicts a reorder buffer (RB) implemented with true FIFO array.

FIG. 20 shows a result-forwarding hazard.

FIG. 21 shows a modified reorder buffer (MRB) according to the invention.

FIG. 22 depicts an MRB circuit implementation.

FIG. 23 illustrates an MRB_cell circuitry.

FIG. 24 shows a support logic for the MRB.

FIG. 25 shows an RBFAQ circuit implementation.

FIG. 26 depicts a support logic for the RBFAQ.

FIG. 27 shows an in-order issue distributed instruction queue (DIQ) according to the invention.

FIG. 28 illustrates an in-order issue DIQ circuit implementation.

FIG. 29 shows a support logic for the DIQ.

FIG. 30 shows an out-of-order issue DIQ according to the 65 invention

FIG. 31 shows a superscalar processor with MRB and DIQ according to the invention.

FIG. 32 depicts a fixed-point register unit (FXRU).

FIG. 33 illustrates a floating-point arithmetic logic unit (FALU).

FIG. 34 shows a fetch and decode e unit (FDU).

5 FIG. 35 shows a flowchart of operations (processes) in fetch and decode stages.

FIG. 36 continues from FIG. 35 and shows a flowchart of operations in issue/execute, writeback, and retire stages.

FIG. 37 depicts a register tag unit (RTU).

FIG. 38 shows a register tag file (RTF).

FIG. 39 shows an RTRB entry.

FIG. 40 shows a register tag reorder buffer (RTRB).

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

The present invention is discussed at length in the abovecited Chamdani's dissertation, the subject of which is incorporated herein by reference.

To achieve the best result in parallelizing a sequential instruction stream, a superscalar microprocessor design typically s supports multi-instruction issue, decoupled dataflow scheduling, out-of-order execution, register renaming, 25 multi-level speculative execution, and precise interrupts. There are two key microarchitecture elements in superscalar hardware that determine the success in achieving the above goal, result shelving and instruction shelving. Result shelving is the key to support register renaming, out-of-order execution, speculative execution, and precise interrupts. Instruction shelving is the key to support multi-instruction issue, decoupled dataflow scheduling, and out-of-order execution.

The following sections discuss, analyze, and compare existing result shelving and instruction shelving techniques. A discussion and analysis of some current state-of-the-art commercial microprocessors are also given. All this information leads to the two superscalar techniques of the invention, the DIQ and MRB, which are described in detail.

Result Shelving Techniques

Result shelving is a technique to temporarily shelve register results before they can be safely retired to a register file. This usually involves buffering of multiple updates to 45 also avoided by copying of operands to the instruction the same register, thus allowing multiple copies/instances of a register to exist in a processor. Four result shelving techniques are discussed in the following sections: register scoreboarding, register-mapping table, checkpoint repair, and reorder buffer. These result shelving techniques are used 50 to support certain superscalar features such as out-of-order execution, register renaming, speculative execution, and precise interrupts.

To support the above superscalar features, the complete processor state, which includes not only register file(s) but 55 also software-visible memory, should be shelved. (It is assumed that a processor's register file(s) contains all architectural registers, including status/control registers.) Stores to the memory can be shelved separately into a store buffer. Each store instruction in the store buffer is committed only 60 if it is safe, meaning: (a) the store instruction can be executed without an exception error (page fault, illegal address), (b) instructions prior to the store instruction are also exception-free, and (c) prior conditional branches have been resolved. These conditions guarantee in-order, non- 65 speculative stores. Thus, despite possible out-of-order and speculative executions inside the processor, the memory is

guaranteed to contain the same state as if the processor executed the code in the original, serial, sequential program order.

Register Scoreboarding

The register scoreboarding technique was originally introduced by Thornton in the CDC 6600 (see J. E. Thornton, Design of a Computer-The Control Data 6600, Scott, Foresman and Co., 1970, M. Johnson, Superscalar Microprocessor Design, Prentice-Hall, 1991). It is the simplest form of 10 result shelving. There is only one level of result shelving, which is accomplished by copying operands and forwarding results to an instruction window. It only allows a maximum of one pending update to a register. Each register in the register file is associated with a scoreboard bit. A "0" 15 indicates the register content is valid, and a "1" indicates there is a pending update. When a register is updated/ written, its scoreboard bit is reset. The term "scoreboarding" used in typical commercial microprocessors, such as the Motorola 88100 and Intel i960CA, does not fit Thornton's definition because they are only used to detect and enforce 20 dependencies (see Johnson, 1991). Instead of buffering to an instruction window, these microprocessors simply stall instruction decoding when data dependencies and resource conflicts are detected.

Thornton's scoreboarding algorithm can be summarized as follows. First, when an instruction is decoded, the scoreboard bit of the instruction's destination register is checked. If it is set, instruction decoding stalls until the destination register is updated by the pending instruction. However, if the destination register's scoreboard bit is clear, instruction decoding proceeds. The decoded instruction is shelved at an instruction window, together with the operand values (if the operand register's scoreboard bit is clear) or the operand register numbers (if the operand register's scoreboard bit is set). Thus, instruction decoding does not stall at data dependencies. The instruction waits at the instruction window until all of its operands become available. Any new instruction result is monitored. If its result register number matches the operand register number, the result value is taken as the 40 operand value. Once all the instruction operands become available, it is ready for issue.

Register scoreboarding is a simple alternative to register renaming. Output dependencies are avoided by allowing only one pending update to a register. Anti dependencies are window. Data dependencies are enforced by forwarding results to the instruction window, to free locked operands.

Register scoreboarding can support out-of-order execution because it preserves data dependencies while avoiding anti and output dependencies. However, it does not support speculative execution nor precise interrupts, because only the most recent updates of registers (architectural state) are maintained, without regard to speculative or out-of-order updates. For speculative execution and precise interrupts to work properly, there must be a mechanism to recover old values and restart at the in-order state prior to the speculative branching or exception instruction.

Register scoreboarding has a serious performance bottleneck. It stalls instruction decoding if there is a pending update to the instruction's destination register. Unlike scalar processors, superscalar processors fetch and decode instructions at a prodigious rate, at two, four, or maybe eight instructions at a time. The probability of an instruction decode stall becomes great. The available instruction parallelism is wasted if there are many register reuses, which is a common practice in most compiler register allocation algorithms.

65

Register-Mapping Table

Keller suggested implementing register renaming by literally renaming the registers (see R. M. Keller, "Look-Ahead Processors," Computing Surveys, Vol. 17, No. 4, pp. 177-195, 1975). A mapping table is used to maintain the mapping or association of an architectural register to its physical register. An architectural register is the register specified in an instruction, which can actually be viewed as the variable name representing a value. A physical register is the true hardware register. The mapping table effectively reintroduces the distinction between register (physical register) and value (architectural register), the essence of register renaming. To avoid anti and output dependencies there must be more physical registers than architectural registers.

The floating-point unit of the IBM RS/6000 (Trademark of IBM, Inc.) uses a 32-entry, 6-bit-wide register mapping table to implement register renaming, as shown in FIG. 2. There are 32 architectural registers and 40 physical registers. Some parts of the register-renaming structure are intentionally omitted from the original diagram to focus on the register mapping table. This register-renaming structure is implemented as a solution to the out-of-order completion problem of floating-point load/store and floating-point arithmetic instructions. In the IBM RS/6000, floating-point load/ store instructions are performed independently at the fixed- 25 point unit, which involve load/store address calculations. Without renaming, out-of-order completion can violate anti or output dependency. For example, in a floating-point load operation, the load data may return too early and overwrite a register that is still needed (i.e., it has not been read by an 30 earlier floating-point arithmetic instruction).

The register renaming process is done as follows. For notational purpose, MP(i)=j (the contents of the mapping table at address i is j) indicates that architectural register Ri maps to physical register Rj. Initially, the mapping table 35 (MP) is reset to identity mapping, MP(i)=i for $i=0, \ldots, 31$. A remapping is performed for every floating-point load/store instruction decoded. Suppose a floating-point load to architectural register 3, FLD R3, arrives at MP. First, the old entry of MP(3), i.e., index 3, is pushed onto the pending-target return queue (PTRQ). Then, a new physical register index from the free list (FL), index 32, is entered to MP(3). This means R3 is now remapped to R32. Further instruction codes with source operand R3 will automatically be changed to R32. Index 3 in the PTRQ is returned to FL (for reuse) 45 only when the last arithmetic or store instruction referencing R3, prior to the FLD R3 instruction, has been performed. This ensures that the current value in physical register R3 is not overwritten while still being used or referenced.

In terms of meeting superscalar objectives, the register-50 mapping table clearly supports register renaming and outof-order execution. However, speculative execution and precise interrupt are not supported because shelving (register mapping) is done only if there are multiple updates to the same register. A speculative or out-of-order update to a register without pending update is not buffered or remapped. 55 With out-of-order execution, recovery from branch mispredictions and exceptions would be impossible. In terms of performance, the register mapping table technique has a disadvantage. The access to the mapping table and the renaming process introduce at least one additional pipeline stage. In the case of the IBM RS/6000, two pipeline stages (PD and Remap) are actually dedicated to the register renaming process.

Reorder Buffer

The reorder buffer (RB) is a content-addressable, circular buffer with head and tail pointers. Entries are pushed in

(allocate) and popped out (retire) in a FIFO (first-in-firstout) manner. It was originally proposed by Smith and Plezkun (see J. E. Smith and A. R. Pleszkun, "Implementation of Precise Interrupts in Pipelined Processors," Proceedings of the 12th Annual Symposium on Computer Architecture, pp. 36-44, 1985) as a hardware scheme to support precise interrupts in pipelined, scalar processors with out-of-order completion. FIG. 3 shows the RB organization. The RB contains the processor's look-ahead state, while the register file (RF) contains the in-order state. The result shift register (RSR) is used to control/reserve the single result bus. (In Smith and Plezkun's processor model, multiple functional units share one result bus.) When an instruction is issued, it reserves stage i of the RSR, where i is the instruction latency (in cycles). If stage i is already 15 reserved, the instruction issue is stalled. The RSR shifts one position every cycle (toward a smaller stage number). When a valid RSR entry reaches stage 1, the result bus control is set such that in the next cycle the result from the entry's functional unit is gated to the RB. 20

There are four operations involving the RB: allocate (entry allocation), read (read operand), write (result write), and retire (entry retirement to RF). To best describe these operations, consider the instruction sequence shown below from (see Smith and Pleszkun).

	Program Example for Reorder Buffer								
	PC	Instructions	Latency						
I0:	0	R2 <- 0	;initialize loop index						
I1:	1	R0 <- 0	initialize loop count						
I2:	2	R5 <- 1	;loop increment value						
I3:	3	R7 <- 100	;maximum loop count						
I4:	4	L1:R1 <- $(R2 + A)$;load A(I)	11 cycles					
I5:	5	R3 < (R2 + B)	;load B(I)	11 cycles					
I6:	6	R4 <- R1 + R3	;floating-point add	6 cycles					
I7:	7	R0 <- R0 + R5	;increment loop count	2 cycles					
I8:	8	(R0 + C) <- R4	;store C(I)						
I9:	9	R2 <- R2 + R5	increment loop index;	2 cycles					
I10:	10	$\mathbf{P} = \mathbf{L}1 : \mathbf{R}0! = \mathbf{R}7$;cond. branch not equal						

Smith and Plezkun state that the RB's allocate operation is performed when an instruction is issued for execution to a functional unit. However, this will restrict the RB to support only out-of-order completion with in-order-issue. To also support out-of-order issue (full or partial), it is necessary to allocate entry when an instruction is decoded. This guarantees that instructions are allocated RB entries in the original program order. For example, FIG. 3 shows the contents of the RB after the RB allocations of I6 and I7. When I7 is decoded, an entry is allocated at the tail of RB (then at entry number 5), following I6's entry. The allocated entry is initialized with "dest. reg."=0 and "program counter"=7. The "valid" bit is also reset to indicate that the R0 value is being computed. The tail pointer is then incremented, modulo RB size. FIG. 3 also shows the RSR contents after I7 is issued (a cycle after I6 is issued). An entry is entered at stage 2 to reserve the result bus, because an integer add takes 2 cycles to complete. The previously allocated RB entry number/tag (5) is written to the RSR entry. The valid bit is set to validate entry.

The RB's read operation is performed prior to instruction issue. Reading an operand directly from the RB (bypassing, without waiting for the RB to retire it to the RF) requires an associative search based on the operand's register number. As an example, suppose I7 is about to be issued. I7's operand register numbers (0 and 5) are presented to both RB

45

55

and RF. The comparators/bypass network compares 0 and 5 to the "dest. reg." field of all RB entries between head and tail pointers. If there is no match then the RF has the latest value. R0 (R5) from the RF is gated to the left_opr_bus (right_opr_bus). If there is a match in RB with "valid"=1, then the "result" field is gated to the left or right operand bus. If the "valid" bit is not set (meaning the operand has not been produced), then I7 is not issued and must wait.

Multi-entry matching can occur in the RB because register number is not a unique associative-search key. Smith 10 and Plezkun refer to this as the "multiple-bypass check" problem, which is also referred to as the "prioritized associative lookup" problem (see Johnson, 1991). obviously, only the "latest" entry should generate a bypass path to the operand bus. For in-order instruction issue, the "latest" entry is simply the last allocated entry among the matched entries. For out-of-order instruction issue, the "latest" entry is the last allocated entry among the matched entries prior to the instruction's RB entry. This becomes the greatest implementation disadvantage for the RB. A significant amount of logic 20 circuitry is required because matched entries can appear in any entry combinations. The RB's write operation is performed when an instruction completes execution. The result value from the result bus is written to the instruction's RB entry ("result" field). The "tag" field in the RSR (stage 1) is 25used to guide the instruction to the correct RB entry. The exception conditions are written to the RB entry's "exceptions" field. To indicate that the result value has been generated, the "valid" bit is set. 30

The RB's retire operation is performed to retire completed result values to the RF (in-order state). Every cycle, the entry at the head of the RB is monitored for retiring. If the result value has already been generated ("valid"=1) and there is no exception ("exceptions"=0), then the head RB entry is retired to the RF. The result value in the "result" field is written based on its register number in "dest. reg.". The head pointer is then incremented, modulo RB size.

In terms of meeting superscalar objectives, it is obvious that the RB supports precise interrupts and out-of-order 40 execution. Results are updated to the in-order state (RF) in the original, sequential program order, after checking their exception conditions. The RB also supports register renaming by creating a new instance/entry for every register assignment. Speculative execution is not readily supported. To support it, requires a mechanism to flush certain RB entries at a variable reset point (branch point). Johnson suggested allocating an entry in the RB for every conditional branch instruction, even though no result value is produced (see Johnson, 1991). When a conditional branch is 50 mispredicted, the processor searches the corresponding "branch" entry in the RB. This entry becomes the reset point; all entries subsequent to it are flushed (tail pointer="branch" entry number). Or alternatively, we can let the "branch" entry reach the bottom (head) of the RB. However, this adds delay to the recovery and restart from a branch misprediction.

Smith and Plezkun presented two other methods for supporting precise interrupts: history buffer and future file. The history buffer does not fit in the result shelving category

because results are written directly to the RF. Any out-oforder issue between two instructions with an identical destination register will cause a wrong register value to be written/read (violate output and anti dependencies). The future file provides a very limited, single-level result shelving and still requires an RB to recover from a branch misprediction or exception. The future file is actually a second RF that is used to store a processor's architectural state, in addition to the first, in-order state RF. Only a single pending update to a register is allowed. Johnson studied these two hardware schemes in supporting speculative execution for superscalar processors (see Johnson, 1991). The history buffer requires excessive cycles to recover and restart from a branch misprediction. The future file also adds 15 delay penalties to mispredicted branches, because the in-order state prior to the branch must be restored to the first RF.

Summary of Comparisons

Table 1 summarizes result shelving techniques, their features and shortcomings. Review of the literature suggests that the reorder buffer (RB) is the most complete result shelving technique (see Table 1). The closest contender to the RB is the register-mapping table, (also called "rename buffer", implemented as a multi-ported lookup table array), which is used in IBM POWERPCs, MIPS R10000 (Trademark of MIPS Technologies, Inc.), and HP PA-8000 (Trademark of Hewlett-Packard, Inc.). The register-mapping table technique has four disadvantages compared to the RB. First, to read a register operand, it has to access the mapping table, using the logical register number, to get the corresponding physical register number in the register file. This additional delay could potentially lengthen the processor's 35 cycle time or introduce another pipeline stage. The third pipeline stage in the MIPS R10000 is dedicated solely to read operand registers.

Second, the mapping table is not a small circuit. For instance, the MIPS R10000 requires two 32×6 mapping tables, one with 12 read ports and 4 write ports for the integer register map, and another with 16 read ports and 4 write ports for the floating-point register map.

Third, precise interrupts are not supported. To overcome this problem, an additional component, called the "active list" in the MIPS R10000, is needed to track every active instruction and maintain the old physical register number prior to the new mapping/renaming (if the instruction updates a register) in the decode stage. If the instruction completes out of order and there is an exception at a preceding instruction, the mapping must be unmapped by writing back the old physical register number from the active list to the mapping table.

Fourth, speculative execution is not supported. To overcome it, the processor must save the whole mapping table in a storage area (called "shadow registers" in the MIPS R10000) for every speculated branch. When a branch is mispredicted, the processor vanks the old mapping table. The storage area increases as more levels of speculative execution are added.

45

50

	Compa	lving Techniques		
	Register Score- board	Register-Mapping Table	Checkpoint Repair (not discussed)	Reorder Buffer
Super- scalar Features Supported	nOut-of-order execution	nOut-of-order execution Register renaming	nOut-of-order execution nRegister renaming nSpeculative execution nPrecise interrupts	nOut-of-order execution nRegister renaming nPrecise interrupts nSpeculative execution (with branch entry location)
Notable Features	ńSimple alternative to register renaming.	nNo associative lookup, direct renaming of register numbers.	nNo associative lookup.	nThe most complete result shelving tech- nique. nFastest branch misprediction recovery and restart
Shortcomi ngs	nMultiple pending updates stall instruction decoding. nDoes not support speculative execution nor precise inter- rupts.	nIntroduces new pipeline stages. nDoes not support speculative execution nor precise interrupts.	nRequires substantial stor- age for the logi- cal spaces. nNeeds complex logic to route result to the appropriate logical spaces. nSlow branch misprediction recovery due to a wait to get the in-order logical space to the Current Space position.	nRequires associ- ative lookup cir- cuitry. nRequires significant amount of comparators and logic circuitry to overcome the pri- oritized associ- ative lookup problem.

Instruction Shelving Techniques

Instruction shelving is a technique to temporarily shelve decoded instructions until they can be issued and executed at the appropriate functional unit. Thus, an instruction shelf (also called instruction window) is a wait station between decode and execution. Instructions in the instruction window can be dynamically scheduled to achieve multiinstruction issue, decoupled dataflow scheduling, and outof-order execution. It is also possible to combine result shelving in the instruction window (called the integrated shelving technique) to support other superscalar features; register renaming, speculative execution, and precise interrupts. An instruction in an instruction window is issued (or ready to be issued) if all of its operand(s) are available, and all resources required (functional unit, operand buses, result bus) are granted.

There are two types of instruction windows: central window and distributed window. In the following sections we will discuss three central window techniques (the dispatch stack, register update unit, DRIS) and one distributed window technique (reservation stations).

Dispatch Stack

The dispatch stack (DS) is a central instruction window that performs dynamic code scheduling on the dynamic instruction stream of multiple functional unit processors. It 65 instructions is decoded and allocated into the window. Since allows out-of-order, multi-instruction issue. The instruction window behaves like a stack where instructions are allocated

at the top and issued from the bottom. After a set of instructions is issued, the gaps (freed entries) are filled with unissued instructions above it (compression). Then, the next set of instructions can be pushed in. This is important to determine instruction order during dependency checks. A DS entry consists of an instruction tag, opcode, source and destination register identifiers, dependence fields, and issue index. To explain how the DS works, consider the program example shown below (see R. D. Acosta, J. Kjelstrup, and H. C. Torng, "An Instruction Issuing Approach to Enhancing Performance in Multiple Functional Unit Processors," IEEE Transactions on Computers, Vol. C-35, pp. 815-828, 1986), which adds floating-point numbers in R0 through R7 and leaves the sum in R0.

55	Program Example with Tight Data Dependencies							
60	I0: I1: I2: I3: I4: I5:	fadd R0,R1,R0 fadd R2,R3,R2 fadd R0,R2,R0 fadd R4,R5,R4 fadd R6,R7,R6 fadd R4,R6,R4	;R0 <- R1 + R0 ;R2 <- R3 + R2 ;R0 <- R2 + R0 ;R4 <- R5 + R4 ;R6 <- R7 + R6 ;R4 <- R6 + R4					
	I6:	fadd R0,R4,R0	;R0 < R4 + R0					

FIG. 4 shows the DS contents after the above sequence of there is no register renaming mechanism in this case, the α and β dependence fields must include not only data

dependencies, but also artificial (anti and output) dependencies. The $\alpha(S1)$ and $\beta(S2)$ count the number of data dependencies (among its preceding instructions in the window). The $\alpha(D)$ counts the number of output dependencies, and $\beta(D)$ counts the number of anti dependencies. An instruction 5 is ready to issue if its issue index (I^2) is zero (meaning no dependencies). In FIG. 4(a), four instructions (IO, II, I3, and I4) are issued, assuming there are adequate functional units and data paths to transmit operands and results. At the completion of each issued instruction, the destination reg- 10 ister identifier (D) is forwarded to the DS, which is compared to the S1, S2, and D fields of all unissued instructions. Wherever there is a match, the appropriate α or β is decremented by 1. FIG. 4(b) shows the updated dependence fields, assuming IO, I1, I3, and I4 complete at the same time. 15 Also note that the window has been compressed. In the next two issues, instructions I2 and I5 are issued simultaneously, followed by I6.

The hardware implementation of a DS is very complex. Each entry requires five counters for the dependence fields 20 and issue index. When an instruction is decoded and allocated to the DS, these counters must be set by comparing the instruction's register identifier with every instruction in the window, requiring (5×number of instructions decoded per cycle×(number of DS entries -1)) comparators. The alloca- 25 tion process also requires each entry to be able to receive an instruction from any decoder position. Since instructions can be issued from any position, every functional unit must be able to receive an instruction from any DS entry. When an issued instruction completes, its destination register identi- 30 fier must be compared to the register identifiers of all instructions in the DS, requiring another set of (5×number of instructions completed per cycle×number of DS entries) comparators. And, what could be the most complex circuitry, receive from any entry above it. Dwyer and Torng estimated that the issue, compression, and allocation circuitry for a mere 8-entry dispatch stack would consume 30,000 gates and 150,000 transistors (see H. Dwyer and H. C. Torng, A Fast Instruction Dispatch Unit for Multiple and Out-of- 40 Sequence Issuances, School of Electrical Engineering Technical Report EE-CEG-87-15, Cornell University, Ithaca, N.Y., 1987 and Johnson, 1991). This complexity makes the dispatch stack unattractive.

Register Update Unit

Sohi and Vajapeyam proposed a simpler central window than the dispatch stack, called the register update unit (RUU) (see G. S. Sohi and S. Vajapeyam, "Instruction Issue Logic for High-Performance Interruptable Pipelined Processors," Proceedings of the 14th Annual Symposium on Computer Architecture, pp. 27-34, 1987). It avoids window compression by keeping issued instructions until they reach the bottom of the RUU, then retire if completed. It has a simpler allocation and issue logic, mainly because it was designed 55 is written back to the instruction's RUU entry at the destifor scalar (single-instruction issue) processors with multiple functional units and long pipeline stages (CRAY-1 category). The RUU resolves data dependencies dynamically, avoids artificial dependencies, and supports out-of-order issue. What is interesting about the RUU design 60 is that it combines the principles of Tomasulo's reservation stations (see R. M. Tomasulo, "An Efficient Algorithm for Exploiting Multiple Arithmetic Units," IBM Journal of Research and Development, Vol. 11, No. 1, pp. 25-33, 1967) and Smith and Plezkun's reorder buffer. Thus, an RUU is an 65 instruction window as well as a result buffer (an integrated shelving technique). It supports five superscalar features;

decoupled dataflow scheduling, out-of-order execution, register renaming, speculative execution, and precise interrupts.

To allow multiple instances of a register (register renaming), each register in the register file is associated with two counters (NI and LI). The NI counter represents the number of instances of a register in the RUU, while the LI counter represents the latest instance number. When an instruction with destination register Ri is decoded and allocated to the RUU, Ri's NI and LI counters are incremented. Wrap around in the LI counter is allowed (modulo counter). However, when the NI counter reaches its maximum value, the instruction decoding is stalled. When an instruction is retired from the RUU and updates the destination register, the NI counter is decremented. With these counters, tag allocations and deallocations become simple. A register tag now simply consists of the register number appended to the LI counter.

In each clock cycle, the RUU performs four operations simultaneously: (a) dispatch/allocate one instruction from the decoder, (b) issue one instruction nearest to the head pointer with ready operands, (c) writeback any result value to the instruction's RUU entry, and forward this result to any matching operand(s), and (d) retire one completed instruction at the head entry and update its result value to the register file. To describe these operations, consider the previous program listing. FIG. 5(a) shows the instruction timing when each instruction is allocated, issued, written back, and retired. It is assumed that each floating-point add takes six cycles to complete. FIG. 5(b) shows the snapshot of the RUU contents at cycle 7. Instruction I6 (fadd R0,R4, R0) has just been allocated at the tail of the RUU. Its program counter, functional unit source, operands, and destination register tag are written. The destination register tag (0,3) is simply the register number (0) appended with the current LI counter value for R0 (3). The "executed" flag is compressing the window requires each entry to be able to 35 reset to indicate entry 6 as unissued. Operands are read directly from the register file. If an operand is available, its value is copied to the allocated entry and the "ready" flag is set. However, if an operand has not been produced, then its register tag is copied to the allocated entry and the "ready" flag is reset. Later when the operand value is produced, the RUU forwards it. By copying operands and forwarding results to the instruction window, anti and output dependencies are effectively eliminated (see Johnson, 1991).

> The issue operation is carried out by first checking 45 "ready" flags of source operand 1 and 2. Any unissued instruction with both operands' "ready" flags set is considered ready. Since only one instruction can be issued per cycle, priority is given to the one nearest to the head pointer. This instruction can issue if all accesses to the functional 50 unit, operand buses, and result bus have been resolved. The issue may be out of order, as shown by the example in FIG. 5(b). Instructions are issued in the following order: I0, I1, I3, I4, I2, I5, and I6.

When an instruction completes execution, its result value nation's "content" field (writeback operation). The "executed" flag is set. To free locked instruction(s) that wait on this result, the destination's "tag" is compared to any source operand's "tag". A match triggers the forwarding of the result value to the operand's "content" field. The "ready" flag is then set. FIG. 5(c) shows the contents of RUU at cycle 9 where RUU performs a writeback operation due to the completion of I1. Because the I1's destination register tag (2,1) matches with I2's second operand, I1's result value is forwarded to I2's source operand 2 "content" and the "ready" flag is set. Now both of I2's operands are available and I2 is ready for issue.

The retire operation is performed on the instruction at the bottom entry of the RUU. If its "executed" flag=1, the result value (destination's "content") is retired to the register file and the head pointer is incremented, modulo RUU size (see I0 in FIG. 5(c)). Retiring from the bottom entry ensures that 5 the register file is updated in the original program order. If the instruction does not complete successfully due to exception or branch misprediction, then a recovery action is initiated by flushing the entire contents of RUU. The register file automatically contains the correct in-order state at the 10 branch point or precise-repair point. The NI and LI counters are reset to zero since the only instance of a register is in the register file (register instances in RUU are completely flushed).

Although the RUU meets many of the superscalar 15 objectives, it is not directly applicable to superscalar processors. The RUU has serious performance bottlenecks (see Johnson, 1991):

- The instruction issue rate is limited to one instruction per cycle (because it was designed for scalar, instead of 20 superscalar processors).
- The window entries of issued instructions are not immediately deallocated. Since only one entry (bottom entry) is freed at a time, a small window size results in frequent stalls at the decoder when the RUU is full. For the RUU to be effective, the window size must be relatively large. Simulation results reported by Sohi and Vajapeyam on Lawrence Livermore loop benchmark programs, show that a relative speedup 30 (compared to the scalar CRAY-1 simulator) greater than 1.8 requires at least 50 entries (see Sohi and Vajapeyam).
- An instruction with branch misprediction or exception is RUU. This is a substantial delay that causes many instructions beyond the branch or exception instruction to be fetched, decoded, and executed unnecessarily, wasting the processor's time with decreased performance.

DRIS

The DRIS (deferred-scheduling, register-renaming instruction shelf) is an improved version of Sohi and Vajapeyam's RUU that is more suitable for superscalar proces- 45 sors. The DRIS is the integrated shelving technique of the Metaf low architecture (see V. Popescu, M. Schultz, J. Spracklen, G. Gibson, B. Lightner, and D. Isaman (Metaflow technologies, Inc.), "The Metaflow Architecture," IEEE DRIS supports all six superscalar features; multi-instruction issue, decoupled dataflow scheduling, out-of-order execution, register renaming, speculative execution, and precise interrupts. The first implementation of DRIS, named DCAF (dataflow content-addressable FIFO), appears in the 55 "locked" bits of unissued instructions ("dispatched"=0). The Metaflow LIGHTNING/THUNDER SPARC microprocessor (Trademark of Metaflow Technologies, Inc.).

The key process to do register renaming, dependency checking, and result forwarding is the tagging of register results. The DRIS uses a different register tagging scheme 60 than RUU. Each register in DRIS is tagged with the identifier (ID) of the producer instruction. The ID of a newly allocated instruction consists of the current "color" bit appended with its DRIS index (entry number). The color bit is used to distinguish the age or order of instructions when 65 the valid entry area (from head pointer to tail pointer) wraps around. When incrementing the tail pointer (modulo DRIS

size), if the DRIS index goes back to zero, the color-bit register is toggled. In determining the order of two instructions X and Y, we compare their IDs. If index(X)>index(Y) and color(X)=color(Y), then instruction X is "younger" (was allocated later) than instruction Y. If their color bits are opposite then the relationship is reversed.

There are four stages/operations in the Metaflow architecture involving the DRIS: dispatch/allocate, issue, writeback, and retire. (To be consistent with the terms used in this document, the Metaflow's terms of "issue", "schedule", and "update" have been changed to the similar terms dispatch/allocate, issue, and writeback, respectively.) To describe these operations, consider the previous program example. FIG. 7(a) shows the instruction timing when each instruction is allocated, issued, written back, and retired. It is assumed that there are four allocate ports, four retire ports, and two floating-point adders with 3-cycle latency. FIGS. 7(b) and 7(c) show the DRIS contents at different time points.

The allocate operation is performed when a set of decoded instructions arrives at the DRIS. For each instruction, the allocated DRIS entry is written with the program counter, opcode, functional unit class number, register numbers of source and destination operand(s), register tags of source operand(s), and some control information. FIG. 7(b) shows a snapshot of the DRIS at cycle 2 when I4, I5, and I6 have just been allocated. The "dispatched" and "executed" bits are initially reset. There are three associative searches involved for each entry allocation; during the setting of "latest" bit in the destination section and during the setting of "locked" bit and "ID" fields in both source operand sections (dependency checking). Consider next the allocation of I6. To determine the "latest" bit, the DRIS allocate logic searches and compares all other entries (including the detected very late, after it reaches the bottom of the 35 ones currently being allocated) with the same destination register number as I6 (0). If no match, the "latest" bit is automatically set. In this case, there are two matched entries (I0 and I2). The ID (color, index) comparisons of I0, I2, and I6 show that I6's entry is the youngest. This means I6 is the 40 last instruction in the DRIS that updates R0. So, I6's "latest" bit is set and I2's and I0's "latest" bits are reset. To determine the "locked" bit and "ID" fields of I6's second operand (R4), register number 4 is associatively searched among the destination registers of other older entries. If no match is found, no instructions in the DRIS updates R4 and R4 must be retrieved from the register file. The "locked" bit is cleared and the "ID" field is set to a default value (e.g. the instruction's own ID) to indicate that the operand is in the register file. However, in this case, two matches are found Micro, pp. 10–13 and 63–73, 1991), shown in FIG. 6. The 50 (13 and 15), but 15's entry contains the latest R4. Therefore, its index (5) plus the current color bit (0) becomes the ID (0,5) of I6's second operand. The "locked" bit and "ID" fields of I6's first operand are determined similarly.

> The issue operation is performed by checking the oldest instruction with both operands unlocked ("locked"=0) is given priority for issue, provided the requisite functional unit is available. The DRIS issue logic strives to find as many instructions to issue as possible. FIG. 7(b) shows a case where the DRIS finds two instructions (I0 and I1) to issue. The DRIS issue logic checks the ID field of each source operand to determine the operand location (either in the register file or the DRIS). Since I0's and I1's source operand IDs are the instruction's own ID (default value for register file location), the operand values are fetched from the register file. In the case that an operand is in the DRIS, the operand ID's index part locates the DRIS entry, where

30

the "content" field of the destination section is then routed as the operand value. In any case, the DRIS issue logic passes the opcode, operand values, and ID of each issued instruction to the corresponding functional unit. The "dispatched" bit of the issued entry is then set.

The writeback operation is performed when an instruction completes execution. The instruction's DRIS entry is found by the result ID's index part. The result value is then written to the "content" field of the destination section. The result ID is broadcast and compared to the operand IDs of other valid DRIS entries. A match clears the corresponding "locked" bit. Thus, the result forwarding unlocks any locked operands that are dependent on the result value. FIG. 7(c) shows how I3's and I4's results unlock both operands of I5.

The retire operation is performed in order, starting from 15 the oldest instruction in the DRIS (bottom entry). Multiple instructions may be retired simultaneously depending on the number of retired ports and the eligibility of instructions. An instruction is eligible to retire if (see Popescu, et al.): (a) it has completed successfully ("executed"=1 and no exception 20 error), (b) all older instructions have been retired or are being retired in this clock cycle, and (c) there is an available retire port to write the register "content" to the register file. FIG. 7(c) shows two instructions (I0 and I1) are eligible to retire. 25

The DRIS handles memory-reference instructions (loads and stores) by forcing them to go through the dispatch/ allocate, issue, execute, and writeback operations twice: first to compute the load/store address, and second to load/store the memory data. Load bypassing is allowed.

The DRIS supports multi-level speculative execution by shelving multiple instances of condition code. The condition code value is written to the DRIS entry of a condition code-setting instruction (at the "content" field of the destination section). When a branch instruction is decoded and 35 generated. A tag is associated with each register instance, allocated a DRIS entry, the branch's operand "ID" field is written with the ID of the youngest condition code-setting instruction (associatively searched). When this instruction completes execution, the writeback operation causes the branch operand to be unlocked. As the oldest branch entry 40 is issued, the predicted branch direction (part of the "opcode" field) is compared to the branch operand, i.e., the condition code fetched from DRIS. If the branch was correctly predicted, then the "executed" bit is set and operation continues as usual (the speculative execution is suc- 45 cessful so far). However, if a branch misprediction is detected, then a branch repair action is invoked. To flush all entries following the mispredicted branch, the DRIS tail pointer is set to the branch ID's index.

Compared to the register update unit, the DRIS improves 50 three areas (see Popescu, et al.): (a) register renaming, (b) operand storage and routing, and (c) branch misprediction delay. First, the implementation of register renaming is improved by a more efficient register tagging (ID). Costly counters per register are avoided. The LI and NI counters 55 also limit the number of instances of a register, which could result in more stalling. Second, the storage per entry in the DRIS is less than the RUU because operand values are not copied to DRIS entries. Rather, the operand values are read directly from the DRIS destination section (or the register 60 file) at issue time. This strategy also saves expensive data routing during result forwarding, eliminating paths from result values to operand fields. Third, branch misprediction is repaired as soon as it is issued and executed, not when it is retired at the bottom of the window. 65

Despite these improvements, the DRIS still carries one RUU disadvantage: window entries of issued instructions

are not immediately deallocated. DRIS entries are freed only when instructions are retired, which proceeds in order, from the bottom of the DRIS. This could result in frequent stalls at the decoder as the DRIS gets full, which may occur due to a small window size or relatively long waits due to many data dependencies or unresolved conditional branches. For the DRIS to be effective, the window size must be sufficiently large, which increases hardware cost. Another potential disadvantage of DRIS is long load and store delays. Many cycles could be saved if the load/store address computation is implemented as a separate instruction and scheduled as early as possible.

Reservation Stations

Tomasulo introduced reservation stations in the floatingpoint section of the IBM 360/91 to exploit the multiple execution units. The main objective was to permit simultaneous execution of independent instructions while preserving the precedence (data dependency) constraints in the instruction stream. Reservation stations are essentially an implementation of distributed instruction windows with some result shelving. The result shelving, done by copying operands and result forwarding to reservation stations, are necessary to support register renaming. FIG. 8 shows the hardware implementation of Tomasulo's algorithm. (To focus the discussion on the reservation stations, two hardware units (floating point buffers and store data buffers) are intentionally omitted from the original diagram.)

There are four key components in Tomasulo's concept: busy bit, tag, reservation station, and common data bus. A busy bit is associated with each floating-point register or operand as a dependency mechanism. If set (busy bit=1) then it means the register is not available, currently being which is used in place of the register number/identifier. This reintroduces the distinction between register and value, the essence of register renaming. In Tomasulo' algorithm, a tag corresponds directly (1-to-1to a reservation station. For example, in the IBM 360/91, tag numbers 10 through 12 correspond to the three reservation stations of the adder unit. However, Weiss and Smith suggested a more flexible way of assigning tags (S. Weiss and J. E. Smith, "Instruction Issue Logic in Pipelined Supercomputers," IEEE Transactions on Computers, Vol. C-33, No. 11, pp. 1013-1022, 1984). When an instruction is decoded, a new tag is assigned for its destination register from a "tag pool" that consists of some finite set of tags. When the instruction completes, the tag is returned to the pool for reuse.

The reservation station is a wait station for an instruction that is about to be executed. Each execution unit has its own set of reservation stations. Each reservation station consists of four fields: sink tag, sink value (first operand), source tag, source value (second operand), and control. The control field contains opcode and destination register number information. It also contains information defining whether the reservation station is ready to issue or not. Two busy bits for the sink and source can be used. If both clear then both operands are available and the reservation station is ready to issue. Note that the IBM 360/91 has two-operand instructions. A sink register is a destination register as well as a source operand. A three-operand processor would have a different reservation station format: source 1's busy bit, tag, value, source 2's busy bit, tag, value, dest tag, and control.

The actions taken during instruction decode are as follows. The decoder decodes one instruction from the top of the floating point operation stack (FLOS). A reservation

60

65

station is allocated at the appropriate execution unit. Instruction operands (sink and source) are copied from FLR to the reservation stations. If the busy bit of an operand register in FLR is clear (indicating the register value is valid), then the register content, tag, and busy bit are copied to the reservation station (at sink or source fields). A new tag is updated to the destination register in FLR, and the busy bit is set. This new tag is the reservation station's assigned number. However, if the busy bit of the operand register in FLR is already set (indicating another instruction is currently generating the register value), then only the register tag and busy bit are copied to the reservation station.

The actions taken during instruction issue are as follows. For each execution unit, one waiting instruction from its reservation station set is selected for issue/execution. This 15 instruction must satisfy the following requirements: (a) all of its operands are available, (b) it has priority over other waiting instructions that are also ready, and (c) its request for the result bus (common data bus), at the clock period when the result will be generated, has been granted. The reserva- 20 tion station number of the issued instruction becomes the destination register tag (dest_tag). The issued instruction then releases the reservation station for others to use.

When the instruction completes execution, the result value, destination register number (dest_reg) and tag (dest_ tag) are placed at the common data bus (CDB). This result is updated to the FLR based on dest_reg and dest_tag. The FLR compares the dest_tag to the last tag of the dest_reg register. The result is also forwarded to the reservation 30 bility (machine parallelism) than the RS. In the RS stations. Each active reservation station compares its sink and source tags to the dest_tag. If they match, the result is written to the sink and/or source fields. This may free some waiting instruction(s).

Implementing instruction shelving by Tomasulo's reser-35 vation stations accomplishes the following superscalar objectives:

- Achieves multi-instruction issue by allowing each funcreservation station set,
- Achieves decoupled dataflow scheduling by shelving all decoded instructions and not stalling instructions at the decoder.
- Permits out-of-order execution while preserving data 45 dependency constraints, and
- Achieves register renaming (enforces data dependencies, eliminates anti and output dependencies) by tagging registers, copying operands and result forwarding to reservation stations.
- Anti dependencies (write-after-read hazards) are avoided by copying operands to reservation stations. Once operands are read from the FLR and copied to the reservation stations (at decode stage), they cannot be overwritten by writes/updates from subsequent instruc- 55 tions. If an operand value is not available from the FLR at the decode stage, its register tag is copied instead to the reservation station. When the operand value is finally generated by one of the execution units, it will be delivered to the reservation station. Thus, result forwarding enforces data dependencies (read-afterwrite hazards). Output dependencies (write-after-write hazards) are avoided by comparing tags at the FLR on every register write, to ensure that only the most recent instruction changes the register.

Tomasulo's algorithm lacks a mechanism to handle speculative execution. If instructions are allowed to be

fetched, decoded, and executed speculatively, then the FLR only contains the architectural state. Only the most recent updates of registers are maintained, regardless of whether they are speculative updates or not. For speculative execution to work properly, there must be a mechanism to recover certain old values and restart at the in-order state prior to the speculative branching. To support speculative execution, the register file can be accompanied by a reorder buffer. The original Tomasulo's algorithm also lacks multi-instruction fetch and decode. A single-instruction decoder really underutilizes the potential machine parallelism. It is not difficult however, to expand Tomasulo's reservation stations to handle multi-instruction decode and become a superscalar distributed window (see Johnson, 1991).

Summary of Comparisons

Table 2 is a summary of instruction shelving techniques. The features which support superscalar design are defined and the good and poor features of each technique are also noted. Review of the literature suggests that the reservation station (RS) is the best instruction shelving technique to give maximum machine parallelism (see Table 2). The closest contender to the RS is the DRIS (Deferred-scheduling, Register-renaming Instruction Shelf), a central window technique used in the Metaflow Thunder SPARC. The DRIS central window technique has five disadvantages compared to the RS distributed window. First, the DRIS has long load and store delays because it forces all memory-reference instructions to go through the dispatch, issue, execute, and writeback pipeline stages twice; first to compute the load/ store address, and second to load/store the memory data.

Second, inherently, the DRIS has less parallel issue capaunit, the number of instructions that can be issued in parallel is as many as the total number of execution units, because each execution unit's window can issue an instruction. On the contrary, the maximum number of parallel issues in the DRIS central window is limited to the number of issue ports.

Third, the hardware to support multiple out-of-order issues is very complex because the issue logic has to deal with greater numbers of instructions and resource allocation conflicts relative to the RS technique, which deals with tional unit to independently issue an instruction from its 40 fewer instructions locally and only needs to issue one instruction.

> Fourth, the instruction routing in the DRIS central window with out-of-order issue requires that there is an issue path from any window entry to any functional unit. The requirement in the reservation-station distributed window is simpler, one issue path from any window entry to one functional unit. Moreover, since the RS window has considerably fewer entries than the DRIS window, the instruction routing complexity becomes even simpler.

Fifth, in DRIS, window entries of issued instructions are 50 not immediately retired/freed, resulting in inefficient use of the window especially if there are long chains of data dependencies or long, unresolved conditional branches. This inefficiency could result in frequent decode stalls and reduce performance. As a result, the DRIS window size must be made relatively large, which increases hardware cost considering each DRIS entry contains many comparators, register bits, and logic to update different kinds of flags (locked, latest, dispatched, executed). The other penalty is the need for DRIS to retire/deallocate multiple entries per cycle. In contrast, the RS only issues one instruction at a time, in which the issued-entry is deallocated immediately.

The proof of DRIS hardware complexity is seen in the Metaflow Thunder SPARC which requires a three-chip set implementation even using IBM's advanced 0.5-micron CMOS process (see R. Wilson, "RISC Face-Off Coming to Hot Chips," Electronic Engineering Times, pp. 1-68, 1994). In general, a central window technique is more complex to implement than a distributed window technique.

TABLE 2

Comparisons of Instruction Shelving Techniques

	Dispatch Stack	Register Update Unit	DRIS	Reservation Stations
Туре	Central window	Central window	Central window	Distributed win-
Supersca- lar Features Supported	nMulti-instruc- tion issue nDecoupled data- flow scheduling nOut-of-order execution	nDecoupled data- flow scheduling nOut-of-order execution nRegister rena- ming nSpeculative ex- ecution nPrecise inter- rupts	nMulti-instruc- tion issue nDecoupled data- flow scheduling nOut-of-order execution nRegister rena- ming nSpeculative execution nPrecise inter-	dows nMulti-instruc- tion issue nDecoupled data- flow scheduling nOut-of-order execution nRegister rena- ming
Notable Features	nImmediately frees issued entries for reuse.	nIntegrated shelving (instruc- tion + result), combining the principles of re- ervation stations and reorder buf- fer. nSimple tag allocation and de- allocation using the NI and LI counters.	Integrated shel-ving (instruction + result), an im- proved version of RUU nNo need to copy operands to DRIS entries, reduces storage require- ment. nSupports multi- level speculative execution.	nIntegrated shel-ving instruction + result). nThe result shel-ving supports register renaming. nMore parallel issue capability than a central window; each func-tional unit has a window to issue
Shortcom- ings	nArtificial dependencies stall instruction issue. nRequires very complex circuitry especially for window compression). nUses many coun- ters and compar- tors. n8-entry dis- patch stack al- ready consumes and 30,000 gates 150,000 transis- tors.	nSingle-instruc- tion issue. nTo be effective, the window size must be relatively large because is- sued entries are not immediately freed. nLate detection of branch mispre- diction and excep- tion. nCopying operands to RUU increases storage/area. Result forwar- ding to all window entries increases global bus rou- ting.	nTo be effective, the window size must be relatively large because is- sued entries are not immediately freed. nLong load and store delays be- cause they are processed in DRIS twice. nThe hardware to support out-of- order is complex, lengthens clock period, and adds a pipeline stage. nResult forwars ding to all DRIS entries increases global bus routing.	instruction. nResult shelving at reservation stations does not support specula- tive execution, unless the regi- ster file is ac- companied by a reorder buffer. nCopying oper- ands to reserva- tion stations in- creases storage and global bus routing. nResult forwar- ding to all reser-vation stations increases global bus routing.

Commercial Superscalar Microprocessors

The first commercial superscalar microprocessor was the 55 Intel i960 CA, which was introduced in 1989. The superscalar features supported were limited, excluding speculative execution and register renaming. The i960 CA decodes four instructions at once, issues up to three of them per cycle, and can sustain a maximum execution rate of two 60 instructions per cycle. The Intel i960 CA was followed by many other versions (CF, MM, etc.). A more complete superscalar processor that supports (single-level) speculative execution and register renaming, the IBM RS/6000, was also introduced in1989 (see R. D. Groves and R. R, Oehler, 65 "An IBM Second Generation RISC Processor Architecture," *Proceedings of the* 1989 *IEEE International Conference on*

Computer Design, pp. 134–137, 1989). Many other superscalar microprocessors followed later.

The following sections describe six commercial superscalar microprocessor architectures: the IBM RS/6000, Metaflow LIGHTNING/THUNDER SPARC, TI SUPER-SPARC (Trademark of Texas Instruments, Inc.), DECchip 21064 (ALPHA AXP), HP PA-7100 (Trademark of Hewlett-Packard, Inc.), and Intel PENTIUM. The discussions focus on the superscalar aspects of the processors. A summary of upcoming superscalar microprocessors that have been recently announced to reach their sampling status in 1995 or 1996 is also given at the end. These include the DEC ALPHA 21164, Sun Microsystems ULTRA-SPARC (Trademark of Sun Microsystems Computer Corporation), IBM POWERPC 620 (Trademark of IBM, Inc.), MIPS

R10000, HP PA-8000, and AMD K5. All of these new microprocessors adopt some kind of superscalar techniques to support out-of-order and speculative execution. Other trends are the migration to true 64-bit architectures, larger bandwidths, multimedia hardware support, and glueless two- or four-way symmetric multiprocessing. Among all existing microprocessors to date, the DEC Alpha 21164, the follow-on to Alpha 21064, currently has the highest performance. Its single-chip architecture is a combination of superpipelined and superscalar architectures. Volume-wise, the Intel PENTIUM is currently the superscalar microprocessor sold in the largest quantities, used in PCs and low-priced workstations, mainly due to its large ×86 software-base.

IBM RS/6000

The IBM RS/6000 is a multi-chip superscalar processor with a RISC instruction set (derivation of the 801 instruction set), capable of executing up to four instructions per cycle. FIG. 9 shows the architecture of the IBM RS/6000. There are three functional units—the branch unit (BU), fixed-point unit (FXU), and floating-point unit (FPU)-that are capable of executing instructions in parallel. The BU can issue up to two instructions per cycle, a branch and a condition-register instruction. The FXU issues one instruction per cycle, which floating-point load/store. The FPU issues one floating-point arithmetic instruction per cycle including a multiply-add instruction. Each functional unit has instruction buffers (I-buffers) to shelve instructions. These I-buffers are orga-BU's I-buffers are a central window that contain all fetched instructions (not decoded yet). The FXU's and FPU's I-buffers are distributed windows. Both receive the same fixed- and floating-point instructions (not decoded at this point).

The instruction cache (I-cache) is a two-way setassociative cache with 16-instruction cache lines. The I-cache can supply four instructions per cycle for each fetch address (PC) presented, regardless of whether the PC is by implementing the I-cache with four independent I-cache arrays, each fetching one instruction per cycle. By adjusting the address and rerouting (aligning) the output of each I-cache array, four instructions can always be fetched, provided they reside in the same cache line. It is to be 45 held off in the pipeline until the branch condition is known. understood that memory or other storage units may be used instead of an I-cache array.

The BU receives four instructions per cycle from the I-cache arrays into the BU's 12-entry I-buffers and dispatch unit. The dispatch unit dispatches externally to the FXU 50 multiply-add-fused unit (MAF), a floating-point register file, and/or FPU any two-instruction combination of fixed- and floating-point instructions. If the remaining two instructions contain a branch and/or a condition-register instruction, they are dispatched and executed internally in the BU. When a conditional branch instruction is encountered, the BU 55 tion (A×B)+C, which reduces the latency for chained fetches the branch-not-taken path instructions (default branch prediction direction) and dispatches them to the FXU and FPU. These instructions are executed speculatively and can be canceled if the branch is mispredicted (by postponing retirement to register files and flushing the pipelines). The 60 branch-taken path instructions are also fetched from the I-cache arrays and placed at the BU's I-buffers, but their dispatching (or flushing) is held off until the branch outcome is known. The worst-case penalty for a branch misprediction is three cycles. The penalty can be eliminated if there are 65 enough independent instructions to separate the compare from the branch.

The BU has four special purpose registers (see R. R. Oehler and R. D. Groves, "The IBM RISC/6000 Processor Architecture," IBM Journal of Research and Development, Vol. 34, No. 1, pp. 23-36, 1990); the machine-state register (to control system states), the link register (for subroutine return address), the count register (to control loop iteration), and the condition register (to support multiple condition codes for conditional branches). Zero-delay branch for loops with a known iteration count is achieved with the branchand-count instruction that uses the count register. The condition register contains eight condition fields, two of which are reserved to contain the condition code results of arithmetic computations in the FXU and FPU. The remaining six can be explicitly set by other fixed- or floating-point com-15 pare instructions and special branch-unit instructions. The setting of each condition field is controlled by the record bit (Rc) of arithmetic instructions. There are advantages to having multiple, setable condition fields. First, the compiler can schedule a compare instruction early, as far away as possible from the conditional branch instruction. Second, 20 several compare instructions can be scheduled first (their results written into separate condition fields), which are then followed by a single conditional branch instruction. This is useful to implement a guarded statement/code section with can be a fixed-point arithmetic, a fixed-point load/store, or a 25 multiple guard conditions, eliminating the typical structure of a series of single compares followed by a single branch.

The FXU contains I-buffers, an arithmetic logic unit (ALU), a general-purpose fixed-point register file, and a single-entry store buffer. The I-buffers receive both fixednized as a FIFO instruction window with in-order issue. The 30 and floating-point instructions from the BU's dispatch unit, but issue only fixed-point instructions and floating-point load/store instructions to the ALU. Addresses of all loads/ stores are computed in the FXU. The ALU includes a multiply/divide unit with 3- to 5-cycle multiply and 19- to 35 20-cycle divide latencies. The store buffer holds data and address of one fixed-point store instruction. The store buffer makes load bypassing possible. Address and data of floating-point store instructions are buffered in the FPU. The I-buffer is a strictly FIFO instruction window with in-order quad-word aligned or not. Instruction aligning is achieved 40 issue. Partial out-of-order issue is achieved among different functional units. Since there is only one ALU and no out-of-order issue in the FXU, the integer RF is not accompanied by a result buffer. Result values are written directly to the integer RF, except on speculative results which are Further instruction issue/execution in the FXU must be stalled. This limits the speculative execution capability in the IBM RS/6000.

> The FPU contains I-buffers, a unified floating-point a register-mapping table, and a store buffer. FPU's I-buffers receive the same instructions as FXU's I-buffers, but issue only floating-point arithmetic instructions to the MAF. The MAF can perform an indivisible multiply-accumulate operamultiply-add operations, rounding errors, chip busing, and the number of adders/normalizers. The latency of a floatingpoint multiply-add instruction (FMA) is two cycles (see R. K. Montoye, et al., "Design of the IBM RISC System/6000 Floating-Point Execution Unit," IBM Journal of Research and Development, Vol. 34, No. 1, pp. 59-70, 1990). The register-mapping table provides register renaming (8 renaming registers) to allow independent, out-of-order executions of floating-point load/store and arithmetic instructions. The store buffer contains five entries for addresses and four entries for data. A floating-point store instruction is issued at the FXU where the store address is calculated and placed at

the store buffer. Once the FPU produces the store value, it is placed at the corresponding entry in the store buffer, ready to be committed to the data cache (it is to be understood that memory or other storage units may be used instead of a data cache). By buffering stores, the FXU can continue issuing 5 subsequent loads (load bypassing).

The number of pipeline stages in the BU, FXU, and FPU are two, four, and six, respectively. To describe how instructions are executed in the RS/6000 pipeline stages, consider the following 2-D graphics transform program example shown below in RS/6000 pseudo-assembly code (see G. F. Grohoski, "Machine Organization of the IBM RISC System/ 6000 Processor," *IBM Journal of Research and Development*, Vol. 34, No. 1, pp. 37–58, 1990).

The IBM RS/6000 processor chipset consists of nine chips (including I-cache and D-cache arrays, bus controller, and I/O controller), which are implemented in a 1- μ m, three-metal CMOS process. The total number of transistors is 6.9 million. The benchmark performance figures on the top of the line system, the IBM RS/6000 POWERSTATION (Trademark of IBM Corporation) 580 (62.5 MHz), are SPECint92 61.7, SPECfp92 133.2 (see R. Myrvaagnes, "Beyond Workstations," *Electronic Products*, pp. 17–18, 1993), and 38.1 MFLOPS on (double precision, N=100) LINPACK (Trademark of Digital Equipment Corporation—see Digital Equipment Corporation, ALPHA AXP Workstation Family Performance Brief-Open VMS, 1992).

In general, the IBM RS/6000 supports all six superscalar features (multi-instruction issue, decoupled dataflow

2-D Graphics Transform Program Example					
I1:	FL	FR0,sin_theta	;load		
I2:	\mathbf{FL}	FR1,-sin_theta	rotation matrix		
I3:	\mathbf{FL}	FR2,cos_theta	constants		
I4:	\mathbf{FL}	FR3,xdis	load x and y		
I5:	\mathbf{FL}	FR4, ydis	;displacements		
I6:	MTCTR	i	load count register with loop count		
I7:_LOOP:	UFL	FR8,x(i)	;load x(i)		
I8:	FMA	FR10,FR8,FR2,FR3	; form $x(i)^*\cos + xdis$		
I9:	UFL	FR9,y(i)	;load y(i)		
I10:	FMA	FR11,FR9,FR2,FR4	; form $y(i)^*\cos + ydis$		
I11:	FMA	FR12,FR9,FR1,FR10	; form $y(i)$ *sin + FR10		
I12:	FST	FR12,x(i)'	store x(i)		
I13:	FMA	FR13,FR8,FR0,FR11	; form $x(i)$ *sin + FR11		
I14:	FST	FR13v(i)'	store v(i)		
I15:	BCT	_LOOP	;continue for all points		

FIG. 10 shows the cycle-by-cycle execution of the inner loop. The superscripts indicate the iteration numbers. During ³⁵ cycle 1, four instructions (17°, 18°, 19°, 110°) starting from_ LOOP label are fetched from the I-cache arrays and placed into BU's I-buffers. During cycle 2, the first load and multiply-add instructions $(17^{\circ}, 18^{\circ})$ are sent to the FXU and FPU, respectively. The next four instructions are fetched 40 (I11^o, I12^o, I13^o, I14^o)). During cycle 3, the FXU decodes the floating-point load (17°) and discards the floating-point multiply-add (18°). The FPU pre-decodes both instructions for register renaming. The loop-closing BCT instruction $(I15^{\circ})$ is fetched. During cycle 4, there is no valid instruction 45 fetch because the branch target address is still being computed. The FXU computes the address of the first load $(I7^{\circ})$, while decoding the second load (I9°). The FPU renames the floating-point registers of $I7^{\circ}$ and $I8^{\circ}$. The BU detects the BCT instruction and generates the branch target address. 50 During cycle 5, instructions from the next iteration $(I7^1, I8^1,$ $I9^1$, $I10^1$) are fetched. The D-cache is accessed for the first load (I7^o). The FXU computes the address of the second load (19°) . The first FMA instruction (18°) is decoded at the FPU. During cycle 6, the FPU executes the first FMA instruction 55 while decoding the second FMA instruction ($I10^{\circ}$). The D-cache is read for the second load (19⁰). In summary, the first iteration outputs x(i)' and y(i)' are stored at cycle 10 and 11, respectively. The iteration period of this loop is 4 cycles. In FIG. 10, there is no branch penalty (zero-cycle branch) in 60 FXU's and FPU's pipelines. The execute pipeline stages (FXE, FPE1, FPE2) are always full, primarily because the instruction fetch rate is twice the instruction issue rate at the arithmetic units. However, a true branch penalty should be seen at the fetch stage (IF), which in this case shows a 65 one-cycle branch delay due to the branch address calculation.

scheduling, out-of-order execution, register renaming, speculative execution, and precise interrupts), some in a limited way. Although four instructions are fetched per cycle, only two (fixed- and floating-point) or three (including a branch) instructions are typically issued per cycle. Only single-level speculative execution is supported. Multiple unresolved conditional branches cause issue stalls because of the lack of a result buffer. Precise interrupts are not supported in the regular mode. They are only supported when the processor is put in the "synchronize" mode, which slows the processor significantly.

Metaflow LIGHTNING/THUNDER SPARC

The LIGHTNING SPARC microprocessor, from Metaf low Technologies, Inc., is the first implementation of the Metaf low architecture that executes the SPARC (v.8) RISC instruction set. The architecture is based on the DCAF (dataflow content-addressable FIFO), a DRIS implementation. Although the DRIS is conceptually a central window, it is implemented as three physical windows: (a) the central DCAF in DIU which shelves all instructions (complete DRIS), (b) the floating-point DCAF in FPU which shelves only floating-point instructions, and (c) the branch DCAF in IIU which shelves only conditional branch instructions. The central DCAF is the central window that is responsible for retiring operations; while others are only a subset of the central DCAF. FIG. 11 shows the LIGHTNING SPARC module which consists of an external cache (up to 1 Mbyte) and four ASICs; the instruction issue unit (IIU), the dataflow integer unit (DIU), the floating-point unit (FPU), and the cache controller/MMU/bus interface (CMB). The external cache consists of the first-level cache for data and the second-level cache for instructions (the first-level is in the IIU chip).

The IIU fetches four instructions per cycle from the on-chip instruction cache (I-cache), assigns a unique ID to each instruction, and computes their PCs. Instruction aligning is achieved by the self-aligning I-cache with an 8-instruction cache line. To guarantee supplying four 5 instructions at all times, a branch to the last three instructions of a cache line causes a read and concatenation of two cache lines. of these four instructions, up to three of them are sent to both the DIU and FPU. If one of the four instructions is a control-transfer instruction (branch, call, return), the IIU 10 executes that instruction itself (at the branch unit). Thus for a typical cycle, all four instructions fetched are not taken/ processed. The IIU also contains a one-entry shelf for the processor state registers. A second instruction attempting to write the same state register before the first has retired, is 15 stalled.

The IIU's branch unit executes a conditional branch instruction speculatively, and shelves it at the branch DCAF for a later misprediction check. The branch DCAF is a specialized DRIS implementation that shelves conditional branch instructions that were speculatively executed. The ID 20 of the oldest, unresolved conditional branch instruction is sent to the retire logic in the central DCAF, to prevent retiring speculative instructions. During the writeback stage, the IDs and results of up to three condition code-setting instructions (two from the integer ALUs in DIU and one 25 from the floating-point adder in FPU) are updated to the branch DCAF. During the execute stage, all branch entries with unlocked operands compare their condition code result with the prediction. If not the same, then the conditional branch was misspeculated. Branch repair is initiated on the 30 oldest mispredicted branch. Its ID is broadcast to all DCAFs. All entries with younger IDs are flushed by moving the tail pointers accordingly.

The DIU contains the central DCAF, two integer ALUs, one memory-address ALU, and retire logic. Up to three 35 instructions can be allocated into the central DCAF, including floating-point instructions which are used for retirement purposes. The central DCAF invokes issue and writeback operations only on integer and memory-reference instructions. To allow proper retiring, each cycle the central DCAF 40 is informed by the FPU on all IDs of successfully executed floating-point instructions. The retire logic can retire up to eight instructions per cycle in the central DCAF (see Popescu, et al.): three that update integer registers or condition codes, two that update floating-point registers, one 45 from Ross Technology, Inc. and Cypress Semiconductor store instruction, and any two instructions of other types (control transfers, processor state updates). The FPU contains the floating-point DCAF, a floating-point adder, and a floating-point multiplier. The floating-point DCAF invokes allocate, issue, and writeback only on floating-point instruct- 50 instruction unit, integer unit, and floating-point unit. There tions. To deallocate entries, the floating-point DCAF is informed by the DIU on the IDs of retired floating-point instructions.

The simulation study shows that the Lightning SPARC can exploit instruction parallelism by mainly relying on its 55 supports up to 1-Mbyte of external cache. dynamic hardware scheduling. The processor is observed to be insensitive to the quality of the SPARC compiler's code generator, especially toward code scheduling and register allocation (see Popescu, et al.). Speculative execution also plays an important role. The processor always has to execute 60 some instructions speculatively, with at least one unresolved conditional branch. The performance limitations observed are memory bandwidth, FPU bandwidth, branch prediction accuracy, branch address calculation latency, and instruction cache hit rate (see Popescu, et al.). 65

The Lightning SPARC never became a commercial product due to the disintegration of the partnership with the original foundry LSI Logic Corp. These shortcomings found during the first design implementation may have also contributed to the failure:

- Low Clock Speed: The out-of-order issue hardware greatly affects control logic delays. Within each clock cycle, the scheduler must examine a large number of instructions in the DCAF, determine which instructions are ready to issue, and select based on priority rules (see Popescu, et al.). This lengthens the basic clock cycle when compared to conventional RISC processors. The estimated clock speed is relatively low, in the 40-MHz range. The out-of-order issue in DCAF also adds a pipeline stage for scheduling.
- No Breakthrough in Performance: The performance on 100×100 Linpack is around 18 MFLOPS while its competitors (the IBM RS/6000, HP PA-7100, DEC ALPHA AXP) are in the 30-40 MFLOPS range. The performance on Dhrystone (Version 1.1) is estimated at 116,000 D/s. The single-chip, 40-MHz SUPERSPARC (on the SPARCSTATION (Trademark of SPARC International, Inc.) 10/41 system) has a Dhrystone performance of 192,400 D/s.
- Expensive Implementation Cost: The design complexity of the DCAF requires substantial VLSI real estate and forces an implementation with four ASICs. Several other single-chip processors with simpler designs have better performance. The medium performance of the Lightning SPARC does not justify the expensive hardware cost and complex design.

In 1994, Metaflow Technologies Inc.'s new partnerships with VLSI Technology Inc. and IBM Microelectronics gave a new birth to the Lightning SPARC design as the Thunder SPARC (see Wilson). With IBM's 0.5-micron, four-metal CMOS implementation the Thunder SPARC chip set is expected to operate at 80 MHz and deliver performance close to 200 SPECint92 and 350 SPECfp92.

TI SUPERSPARC

The SUPERSPARC processor from Texas Instruments, Inc. is the first commercial superscalar implementation of the SPARC version 8 architecture (Sum Microsystems Computer Corporation, The SUPERSPARC Microprocessor-Technical White Paper, 1992). A virtually identical version Corporation is called the HYPERSPARC (Trademark of Ross Technology, Inc. and Cypress Semiconductor Corporation). FIG. 12 shows the SUPERSPARC architecture which primarily consists of three functional units: an are also a 20-Kbyte I-cache that fetches four instructions per cycle, and a 16-Kbyte D-cache that can handle one 64-bit load or store per cycle. These on-chip caches can interact with the MBus or a second-level cache controller that

The instruction unit is responsible for instruction fetch, decode, issue, and branch execution. It fetches four instructions from the I-cache to either the 8-entry sequential instruction queue (for sequential or untaken branch path) or the 4-entry target instruction queue (for taken branch path) (see F. Abu-Nofal, et al., "A Three-Million-Transistor Microprocessor," Digest of Technical Papers of the 1992 IEEE International Solid-State Circuits Conference, pp. 108-109, 1992) The SUPERSPARC always predicts all branches are taken (see Sun Microsystems Computer Corporation, 1992), and fetches instructions in the branchtarget path with one-cycle branch delay (to calculate the

20

target address). The selected instruction queue is essentially a central instruction window with in-order issue. Three oldest/bottom instructions from the selected instruction queue are presented to the scheduling and grouping logic block, where: (a) opcodes are decoded, (b) instruction grouping is determined, (c) intergroup data dependencies are checked, (d) resources are allocated, and (e) bypassing is controlled. Not all three candidate instructions can be issued. There are some grouping restrictions (a total of 23 rules) such as (G. Blanck ans S. Krueger, "The SUPERSPARC 10 Microprocessor," Proceedings of the 37th COMPCON, pp. 136-142, 1992):

Maximum of two integer results,

Maximum of one shift instruction,

Maximum of one floating-point arithmetic instruction,

Maximum of one control transfer (branch) instruction,

Certain "hard" instructions (e.g., subroutine save/restore, integer multiply/divide, control-register update) are issued as a single-instruction group exclusively.

Instructions are never issued out of order. Thus if the third instruction is issued then so are the first and second instructions. Remaining instruction(s) that cannot issue are recirculated to the sequential instruction queue.

The integer unit executes all integer arithmetic instruc- 25 tions (except integer multiply and divide) and load/store address calculations. A virtual address adder and two RF read ports are dedicated to load/store address calculations. Three ALUs are dedicated to arithmetic instructions. The shift instructions can only execute at the first-level ALU with shifter. The second-level ALU is provided for "cascaded" arithmetic operations, to allow back-to-back dependent integer instructions within an instruction group.

The floating-point unit provides a 4-entry floating-point instruction queue, 5-port floating-point register file, floating- 35 point adder, and floating-point multiplier. A floating-point instruction is issued from the bottom (oldest entry) of the instruction queue when the operands and resources are available. All floating-point instructions start in order and complete in order (see Sun Microsystems Computer 40 Corporation, 1992). The floating-point adder performs addition, subtraction, format conversion, comparison, absolute value, and negation. The floating-point multiplier performs single- and double-precision multiplication, division, square root, and integer multiplication and division. Bypass- 45 ing capabilities from the result buses and load bus to arithmetic units are provided. The latency of most floatingpoint operations is three cycles.

The SUPERSPARC processor is implemented using a 0.8 μ m, three-metal BiCMOS process. It integrates 3.1 million 50 transistors and currently runs at 50 MHz. The pipeline is implemented based on a two-phase, non-overlapping clocking scheme. Instruction processing consists of eight pipeline stages/phases: two phases of instruction fetch (F0, F1), three phases of decode (D0, D1, D2), two phases of execution (E0, 55 ports and instruction issue buses. If an instruction pair E1), and a writeback phase (WB). The top benchmark performance figures are: SPECint92 68 and SPECfp92 85.

The SUPERSPARC processor supports five superscalar features; multi-instruction issue, decoupled dataflow scheduling, speculative execution, out-of-order execution, 60 and precise interrupts. Speculative execution is handled by holding a speculative result at the end of the pipeline before being written to the register file (until the conditional branch

is resolved). If mispredicted, all instructions and results currently in processing pipelines are flushed. The multiplepath fetching into the sequential and target queues helps reduce the branch misprediction penalty. The SUPER-SPARC architecture is still somewhat limited in its superscalar capabilities: (1) The multi-instruction issue has a lot of restrictions/rules, and is limited to three instructions despite the four-instruction fetch. Output dependencies also stall instruction issue because register renaming is not supported.

(2) Only single-level speculative execution is supported because of the lack of a result buffer.

(3) Only limited out-of-order execution is supported; no load bypassing, and strictly in-order issue with the possibil-15 ity of out-of-order completion of floating-point instructions from the floating-point queue with respect b integer instructions.

DEC ALPHA 21064

The DEC ALPHA 21064 processor is the first implementation of Digital Equipment Corporation's 64-bit ALPHA AXP architecture (see E. McLellan (Digital Equipment Corporation), "The APLHA AXP Architecture and 21064 Processor," IEEE Micro, pp. 36-47, 1993). It is currently the fastest single-chip microprocessor in the industry. The architecture is a combination of superpipelined and superscalar architectures. The integer and floating-point pipelines are seven- and ten-stages deep, respectively. Since DEC has an existing, large customer base of software, it offers compatibility with VAX and MIPS codes through binary translation. Executable program codes are converted to AXP code without recompilation (with some performance degradation). FIG. 13 shows the DEC ALPHA 21064 architecture, which has four functional units: an instruction unit (IBox), an integer unit (EBox), a floating-point unit (FBox), and an address unit (ABox). There are also 32 entry by 64-bit integer and floating-point register files (RFs), 8-Kbyte D-cache, and 8-Kbyte I-cache with a 2K by 1-bit branch history table. The branch history table is provided for dynamic prediction and achieves 80% accuracy on most programs. Static prediction is also supported based on the sign of the branch address displacement field as the default; backward branches are predicted taken and forward branches are predicted not-taken.

The IBox fetches a pair of instructions from the on-chip I-Cache, decodes them, and issues up to two instructions that pass all register conflict checking. Branch instructions are also executed in this unit. The virtual branch target address is translated to its physical address, which is simply accessed from the I-TLB (instruction translation look-aside buffer). The DEC ALPHA 21064 restricts the instruction pairs that can be issued simultaneously, because of the limited RF cannot dual issue, either the pair is swapped with another pair capable of dual issue, or the pair is serialized, issuing one instruction at a time in order. Although it is not mentioned in McLellan, it is interpreted that pair swapping and serialization operations require some kind of shelving of fetched instructions into a prefetch buffer (central instruction window). Based on the rules shown below, the IBox determines which instruction pairs can dual issue.

Dual Instruction Issue Rules in the DEC ALPHA 21064 (see McLellan)							
Valid Dual Issue:	Instruction A	Instruction B					
	integer operate integer/floating-point load integer store floating-point store integer store floating-point store integer branch floating-point branch	floating-point operate integer/floating-point operate/branch integer operate integer branch floating-point branch floating-point operate integer operate floating-point operate					

Note that a branch instruction is formatted as an integer (floating-point) instruction if its condition code is in an integer (floating-point) register. The DEC ALPHA 21064 avoids condition codes, special registers, or any other single copy of a resource which can potentially become a point of $_{20}$ contention in a multi-instruction issue environment. Compare instructions write directly to any general-purpose register (integer or floating-point, depending on the compare operation type).

The EBox contains dedicated integer multiplier, adder, 25 shifter, and logic units. The multiplier unit is not pipelined to save silicon area. The adder and logic units have singlecycle latency with bypass paths for register write data. The shifter takes two cycles to produce results, but is fully pipelined (one-cycle throughput). The FBox contains dedi- 30 cated floating-point multiplier/adder and divider units. It supports both VAX- and IEEE-standard data types and rounding modes. The divider unit generates one bit of quotient per cycle. All other floating-point operate instructions have six-cycle latency and one-cycle throughput.

The ABox performs all load and store instructions. It has a dedicated displacement adder to compute load/store addresses independently from the IBox. A 4-entry×32-byte store buffer is provided for load bypassing and merging of data from adjacent stores to increase effective bandwidth. A 40 memory barrier instruction is provided to disable load bypassing when necessary. The ABox also contains a 32-entry data TLB to translate the virtual load/store address to its physical address, and 3-entry load silo to buffer outstanding load misses. With a hit at the primary D-cache, 45 the latency of a load is three cycles.

FIG. 14 shows the pipeline stages of the DEC ALPHA 21064 processor for integer and floating-point instructions. Up to two instructions can be processed in each stage. The first three stages (IF, SW, IO) can be stalled, while stages 50 beyond IO advance every cycle (see D. W. Dobberpuhl, et al., "A 200-MHz 64-Bit Dual-Issue CMOS Microprocessor," Digital Technical Journal, Vol. 4, No. 4, Special Issue, pp. 35-50, 1992). In stage IF, a pair of instructions is fetched from the on-chip I-cache. In stage 55 SW, a swap or serialization operation is performed based on the dual-issue rules. If there is a conditional branch instruction, the branch direction is predicted statically or dynamically (using the branch history table). In stage IO, instruction(s) are decoded and checked for dependencies 60 between the two fetched instructions (if any). In stage II, instructions) are issued to the appropriate functional unit, provided there is no register conflict. The source operands are read from the integer and/or floating-point RFs and sent to the EBox, IBox, ABox, and FBox. In stage 4, instruction 65 executions start (stage A1 for integer instructions, stage F1 for floating-point instructions).

FIG. 14(a) shows that all integer arithmetic and logic instructions (EBox), except shift instructions, have onecycle latency, through bypass paths. Shift instructions have two-cycle latency. All results in EBox are actually written back to the integer RF in stage 6. Without the bypass path, the latency would be three cycles. But with the bypass path, the latency is reduced to one or two cycles. This improves the probability that back-to-back dependent instructions execute at full pipeline speed. The DECchip 21064 dedicates 45 different bypass paths. Conditional branch instructions (IBox) are resolved in stage 4. If a branch misprediction is detected, a branch repair is initiated. Instructions subsequent to the branch (in the wrong path) and their intermediate results are flushed from all pipeline stages. The alternate branch target address is computed as the new PC. The first instruction pair of the correct branch path is fetched at stage 6. This branch misprediction causes a four-cycle delay. Primary, on-chip D-cache accesses of load and store instructions (ABox) complete in stage 6. So, the latency of loads and stores is three cycles. FIG. 14(b) shows that results of floating-point operations (from the multiplier/adder unit) are written back to the floating-point RF in stage 9, thus giving a 6-cycle latency. The ALPHA AXP architecture has several notable characteristics:

- Design Simplicity: The architecture avoids direct hardware support of features that are seldom used or prone to limit performance due to cycle-time restrictions. For example, it omits support for direct-byte load/store instructions and precise arithmetic exceptions. Some functionality sacrifices have to be made to get a very high processor clock frequency. Full functionality is achieved through software assistance.
- Privilege Architecture Library (PAL): The architecture supports multiple operating system (O/S) ports (currently OpenVMS, 64-bit Unix DEC OSF/1, Microsoft WINDOWS NT (Trademark of Microsoft Corporation)) using a set of privileged software subroutines, the PALcode. The PLACODE routines are written by the O/S programmer and can be completely customized since they use a superset of the AXP instruction set. They can implement lowest-level hardware-related tasks unique to a particular O/S, such as interrupts, exceptions, context switching, memory management, etc. The PALcode routines are invoked by hardware traps or explicit CALL_PAL instructions.
- Conditional Move Instructions: The AXP instruction set includes conditional move instructions for both integer and floating-point data. These instructions should help remove some conditional branches (see section below entitled Condition Move Transformation).
- Imprecise Interrupts: Precise exceptions are not supported. A user must use the trap barrier instruction

60

when precise interrupt behavior is necessary. In this case, the performance is sacrificed.

The DEC ALPHA 21064 single-chip processor is implemented using a 0.75 μ m, three-metal CMOS process, with operating speeds up to 200 MHz. The extremely high clock 5 frequency presents a difficult clocking situation. To avoid race conditions for latched data, the clock edge rate must be extremely fast (0.5 ns) and only very little clock skew can be tolerated. DEC's solution is to implement a very large, on-chip clock driver with a final stage containing 156 to 10 172-mil-wide pMOS and 63 to 78-mil-wide nMOS devices (see McLellan). The clock driver occupies about 5% of the total chip area and draws a peak switching current of 43 A. A $0.13 - \mu$ F on-chip decoupling capacitance must be added to dissipation is 30 W at 200 MHz with a 3.3-V supply. Sophisticated packaging is used to cool the chip. These hardware cost and implementation problems are compensated by top performance. The benchmark performance figures on the top-of-the-line system, the DEC 10000/610 20 (200 MHz), are: SPECint92 116.5, SPECfp92 193.6, and 40.5 MFLOPS on 100×100 Linpack (double precision).

The DECchip 21064 processor supports three superscalar features; multi-instruction issue (dual issue), decoupled dataflow scheduling, and limited out-of-order execution (by 25 load bypassing and pair swapping for dual issue). The load bypassing permits out-of-order execution between loads and stores. The pair swapping selects the two oldest instructions capable of dual issue, resulting in out-of-order issue. The DEC ALPHA 21064 does not push the superscalar design aggressively since the main goal is to achieve very high clock frequency. Some of the restrictions are:

- (1) Only two instructions are fetched per cycle, which are issued with some pairing restrictions. DEC has just tackled this problem by having a quad-issue on its next 35 architecture. generationALPHA21164.
- (2) Register renaming is not supported which inhibits implementation of full out-of-order execution. Anti and output dependencies will stall instruction issue.
- (3) No speculative execution is supported due to the lack of a result buffer. A conditional branch must be resolved at stage 4 before the predicted path can proceed further into the execution pipeline. An unresolved conditional branch will stall the decode and issue pipelines.
- (4) Precise interrupts are not supported in hardware. The software solution complicates debugging and slows down the processor significantly.

HP PA-7100

The HP PA-7100 processor is the seventh implementation of Hewlett-Packard's PA-RISC (precision architecture, reduced instruction set computer-Trademark of Hewlett-Packard, Inc.) architecture (T. Asprey, et al. (Hewlett-Packard), "Performance Features of the PA-7100 55 (SMU), dedicated branch adder, and a 32×32-bit, general-Microprocessor," IEEE Micro, pp. 22-35, 1993). It is the first superscalar PA-RISC design, which issues up to two instructions per cycle. Its design also has a VLIW flavor. There are two notable design approaches in the PA-RISC architecture; (a) the use of off-chip, rather than on-chip, primary caches (I-cache and D-cache), and (b) the reduction of instruction count in programs (pathlength reduction-see R. Lee, et al., Pathlength Reduction Features in the PA-RISC Architecture," Proceedings of the 37th COMPCON, pp. 129-135, 1992) by adding VLIW-like and SIMD-like 65 instructions. The motivation to use off-chip caches is the fact that on-chip caches are usually not large enough to achieve

balanced performance across a wide range of applications. Typically, on-chip I-caches range from 8 to 20 Kbytes, and on-chip D-caches range from 8 to 16 Kbytes. The PA- 7100 processor can have up to 1 Mbyte I-cache and 2 Mbyte D-cache. Unlike most processors with small on-chip caches, a secondary cache becomes unnecessary. Another advantage is the flexibility of cache size and speed to configure different systems, from low-end to high-end systems.

The objective of pathlength reduction is to resolve the key disadvantage of RISC architectures, the code/pathlength expansion. There are two instruction types added to the RISC instruction set. First, two or three operations that frequently occur together are combined into a fixed-length, 32-bit instruction. This results in multi-operation, VLIWovercome the supply voltage problem. The chip's power 15 like instructions (except they are contained within a short 32-bit instruction), such as Shift&Add (perform integer multiplications with a small constant), Multiply&Add (floating-point), Compare&Branch, Add&Branch, Branch_ on_Bit, etc (see Lee, et al.). Other streamlined RISC architectures such as MIPS require multiple instructions to perform these tasks. Second, SIMD-like instructions are added to operate, in parallel, on multiple data units smaller than a 32-bit word. These instructions are particularly useful in parallel character and decimal operations. For example, in the C language, character manipulations frequently involve finding the null byte (zero) that marks the end of a variablelength string of characters. PA-RISC's Unit_Exclusive_Or instruction speeds this process by testing a "no byte zero" in a word of four bytes in a single cycle (see Lee, et al.). The $_{30}$ addition of the two instruction types is accommodated in the hardware without impacting the cycle time or the CPI. This gives the PA-RISC architecture some of the advantages of a very simple VLIW architecture (with short 32-bit instructions), without losing the advantages of a RISC

> FIG. 15 shows the PA-7100 architecture. The processor chip consists of six major blocks; the integer unit, floatingpoint unit, cache control/interface, unified TLB, control unit, and system bus interface. The control unit is responsible for fetching, decoding, and issuing of instructions. Two instructions are fetched from the off-chip I-cache per cycle, and buffered in a small prefetch buffer (central window). The control unit can issue up to two instructions per cycle, one to the integer unit and one to the floating-point unit. There 45 are no alignment or order constraints on the pair of instructions (see E. DeLano, et al., "A High Speed Superscalar PA-RISC Processor," Proceedings of the 37th COMPCON, pp. 116-121, 1992). However, no two integer or floatingpoint instructions can be issued simultaneously. If a condi-50 tional branch instruction is encountered, a simple static branch prediction scheme is used to minimize branch penalty. All forward conditional branches are untaken and backward conditional branches are taken.

The integer unit contains an ALU, shift-merge unit purpose, integer register file. Besides integer arithmetic instructions, the integer unit also executes branch instructions, loads and stores of integer and floating-point registers, and all VLIW-like and SIMD-like instructions, except the floating-point Multiply&Add and Multiply&Subtract. The VLIW-like instructions improve the utilization of the three hardware units. For example, the Add&Branch uses the ALU and the branch adder simultaneously, while the Branch_on_Bit uses the SMU and the branch adder (see Lee, et al.). The register bypass paths produce a one-cycle latency for integer arithmetic instructions.

The floating-point unit contains a floating-point ALU (FALU), multiplier (FMUL), divide/square root unit (FDIV/ SQRT), and a 32×64-bit floating-point register file. Although there are 32 physical registers, the first four registers (0-3)are dedicated for status register and exception registers. The 5 remaining 28 registers (4-31) are used as register operands for arithmetic operations. Each register can be access as a 64-bit double word or as two 32-bit single words. The FALU performs single- and double-precision add/subtract, compare/complement, and format conversion instructions. 10 The FMUL performs single- and double-precision multiplications, and also 32-bit unsigned integer multiplications (64-bit result). The multiplier array is based on a radix-4 Booth encoding algorithm. The register bypass paths produce a two-cycle latency for all floating-point instruc-15 tions performed in the FALU and FMUL. The FDIV/SQRT performs floating-point divide and square-root operations based on a modified radix-4 SRT (Sweeney, Robertson, and Tocher-see Asprey, et al.) algorithm. The main modification is running the radix-4 division hardware at twice the 20 processor clock frequency to effectively achieve a radix-16 performance. Four quotient bits are computed each clock cycle, giving a latency of 8 and 15 cycles for single- and double-precision divide/square root operations. The floating-point register file has five read ports and three write ports to allow concurrent execution of a floating-point multiply, a floating-point add, and a floating-point load or store. This occurs when a Multiply&Add or a Multiply&Subtract instruction is issued concurrently with a ger instruction).

The instruction execution pipeline for various types of instructions is shown in FIG. 16. The pipeline frequency is determined by the read cycle time of the off-chip cache into two equal phases (2-phase clocking scheme). The first three phases are dedicated for instruction fetching from the off-chip I-cache. Instruction decode and issue can be done in a mere single phase because a pre-decoded bit is dedicated integer and floating-point units. The phases for instruction execution depend on the instruction type, as depicted in FIG. 16. For a conditional branch instruction, instructions along the predicted path are fetched (static branch prediction) Corporation, "Architecture of the PENTIUM while the branch condition is evaluated. In the meantime the $_{45}$ Microprocessor," *IEEE Micro*, pp. 11–21, 1993). alternate link address (Laddr) is also calculated. If at the end of the execute stage the branch is found to be mispredicted, the previous speculative instruction fetch is flushed and new instructions along the correct path are fetched. If the delay is viewed from the I-fetch to the Target I-fetch, the minimum 50 and performs instruction aligning. Because ×86 instructions branch delay of correctly and incorrectly predicted branches is one cycle and two cycles, respectively. The PA-7100 processor has extensive register bypass capability to minimize pipeline interlock penalties. As illustrated in FIG. 16, the penalty for integer ALU pipeline interlock is zero cycles. 55 The penalty for load use, floating-point ALU, or floatingpoint multiply pipeline interlocks is one cycle.

The HP PA-7100 processor is implemented using a 0.8 μ m, three-metal CMOS process. It operates at 100 MHz and integrates about 850,000 transistors. The use of off-chip caches results in a large pin-count package, 504-pin PGA. The reported benchmark performance figures on the top-ofthe-line system, the HP9000/735 (99 MHz), are: SPECint92 80.0, SPECfp92 150.6, and 40.8 MFLOPS on 100×100 Linpack (double precision).

60

65

The HP PA-7100 processor supports four superscalar features; multi-instruction issue, decoupled dataflow

scheduling, out-of-order execution, and precise interrupts. However, it does not support register renaming and speculative execution. Note that static branch prediction is only used for speculative prefetch, not speculative execution. Instructions following an unresolved conditional branch are stalled and not executed. The HP PA-7100 designers rely on aggressive VLIW-like software scheduling and chose not to push the superscalar hardware too aggressively:

- (1) The multi-instruction issue is limited to dual issue of integer and floating-point instruction pairs. No two integer or floating-point instructions can be issued simultaneously. To increase machine parallelism, the VLIW-like and SIMD-like instructions are included, which increases the complexity of the compiler.
- (2) The decoupled dataflow scheduling is limited by a small central window (prefetch buffer) that only issues instructions in order.
- (3) The out-of-order execution is limited to in-order issue with the possibility of out-of-order completion.

Intel PENTIUM

The Intel PENTIUM microprocessor is the first superscalar implementation that runs the widely-used ×86 CISC instruction set. The ×86 instructions use only two operands and permit combinations of register and memory operands. Thus, unlike all other commercial superscalar processors, the PENTIUM processor is not a typical register-to-register, three-address machine. Despite the complexity of CISC floating-point load/store instruction (categorized as an inte- 30 instructions, many of which require microcode sequencing, the PENTIUM processor manages to differentiate the "simple" (RISC-like) instructions and executes them in superscalar mode (dual-instruction issue). However, complex instructions and almost all floating-point instructions RAMs (see Asprey, et al.). Each pipeline stage is divided 35 must still run in scalar mode (single-instruction issue). The superscalar execution and architectural improvements in branch prediction, cache organization, and a fully-pipelined floating-point unit result in a substantial performance improvement over its predecessor, the i486 processor. When in the instruction-field format to steer instructions to the 40 compared with an i486 processor with identical clock frequency, the PENTIUM processor is faster by factors of roughly two and five in integer and floating-point performance, respectively (D. Alpert and D. Avnon-Intel

> FIG. 17 shows the PENTIUM architecture. The core execution units are two integer ALUs and a floating-point unit with dedicated adder, multiplier, and divider. The prefetch buffers fetch a cache line (256 bits) from the I-cache are of variable length, the prefetch buffers hold two cache lines; the line containing the instruction being decoded and the next consecutive line (see Alpert and Avnon). An instruction is decoded and issued to the appropriate functional unit (integer or floating-point) based on the instruction type. Two instructions can be decoded and issued simultaneously if they are "simple" instructions (superscalar execution). Because ×86 instructions typically generate more data memory references than RISC instructions, the D-cache supports dual accesses to provide additional bandwidth and simplify compiler instruction scheduling algorithms (see Alpert and Avnon).

> The integer unit consists of two integer pipelines, the U pipe and V pipe. Two consecutive integer instructions I1 and I2 can be issued simultaneously to U pipe and V pipe, respectively, if: (a) I1 and I2 are from the class of "simple" instructions, (b) I1 is not a jump instruction, (c) destination

of I1 \neq source of I2 (no data dependency), and (d) destination of I1 \neq destination of I2 (no output dependency). Otherwise, only I1 is issued to the U pipe. Note that an instruction issue is stalled on output dependency (artificial dependency) because the PENTIUM processor does not implement register renaming. However, the PENTIUM anti dependencies are of no concern since instructions are issued in order and operand reads occur in an earlier pipeline stage than result writes. The V pipe can execute ALU operation, memory reference, and jump instructions. The U pipe can execute an additional set that uses special hardware available only in the U pipe, such as the barrel shifter. The "simple" instruction class is limited to instructions that can be executed directly and do not require microcode sequencexecuted in the integer SPEC benchmark suite (SPECint92) are "simple" (see Alpert and Avnon).

The PENTIUM processor employs dynamic branch prediction using a branch target buffer (BTB), which holds entries for 256 branches. When a conditional branch instruc- $_{20}$ tion is decoded, the processor searches the BTB to find if there is a history (taken or untaken) on the branch instruction. If the branch has untaken history or no history at all, then the processor continues to fetch and decode instructions along the sequential path. If there is a taken history, then the $_{25}$ processor uses the target address in the BTB entry to fetch and decode instructions along the target path. The conditional branch is resolved early in the writeback stage. If the branch is mispredicted, the processor flushes the pipeline and resumes fetching along the correct path. The branch 30 history in the BTB is updated. This speculative execution with the BTB allows the PENTIUM processor to execute correctly-predicted branches with no delay. Also, because a conditional branch is resolved after the execute stage, a conditional branch instruction in the V pipe can be paired 35 with a compare instruction or other instruction in the U pipe that sets the condition flag.

The floating-point unit consists of six functional blocks: the floating-point interface, register file, and control (FIRC), the floating-point exponent (FEXP), the floating-point mul- 40 tiplier (FMUL), the floating-point adder (FADD), the floating-point divider (FDIV), and the floating-point rounder (FRND). The FIRC contains a floating-point register file, interface logic, and centralized control logic. The ×86 floating-point instructions treat the register file as a stack of 45 eight registers, with the top of the stack (TOS) acting as the accumulator. They typically use one source operand in memory and the TOS register as the other source operand as well as the destination register. In the case of 64-bit memory operands, both ports of the D-cache are used. To swap the 50content of the TOS register with another register, the FXCH instruction (non-arithmetic floating-point instruction) is used. The FIRC also issues floating-point arithmetic instructions to the appropriate arithmetic blocks. Non arithmetic floating-point instructions are executed within the FIRC 55 itself. Floating-point instructions cannot be paired with any other integer or floating-point instructions, except FXCH instructions.

The FEXP calculates the exponent and sign results of all floating-point arithmetic instructions. The FADD executes 60 floating-point add, subtract, compare, BCD (binary coded decimal), and format conversion instructions. The FMUL executes single-, double-, extended-precision (64-bit mantissa) floating-point multiplication and integer multiplication instructions. The FDIV executes floating-point 65 divide, remainder, and square-root instructions. And, the FRND performs the rounding operation of results from

FADD and FDIV. The floating-point unit also supports eight transcendental instructions such as sine (FSIN), cosine (FCOS), tangent (FPTAN), etc. through microcode sequences. These instructions primarily involve the FADD arithmetic block and sometimes other arithmetic blocks.

The PENTIUM processor is implemented using a 0.8 μ m BiCMOS process. It integrates 3.1 million transistors and currently runs at 66 MHz. The integer pipeline consists of five stages: prefetch (PF), first decode (D1), second decode (D2), execute (E), and writeback (WB). The floating-point pipeline consists of eight stages, where the first three stages (FP, D1, and D2) are processed with the resources in the integer pipeline. The other floating-point stages are: operand ing. Intel's study shows that more than 90% of instructions ¹⁵ (WF), and error reporting (ER). The reported benchmark fetch (E), first execute (X1), second execute (X2), write float performance figures of the PENTIUM processor are: SPE-Cint92 64.5 and SPECfp92 56.9.

> The PENTIUM processor supports five superscalar features; multi-instruction issue, decoupled dataflow scheduling, speculative execution, out-of-order execution, and precise interrupts (using the safe instruction recognition mechanism (see Alpert and Avnon). However, most of these features still have serious limitations which are supposed to be improved in the P6 and P7 designs:

- (1) The multi-instruction issue is generally limited to two "simple" integer instructions with certain restrictions. Output dependencies (artificial dependencies) stall instruction issue because register renaming is not supported Floating-point instructions also cannot be paired with any other instructions, except occasionally with FXCH instructions (but FXCH may be considered as a useless or unnecessary instruction in true register-toregister architectures). The multiple floating-point arithmetic blocks (FADD, FMUL, FDIV) are underutilized by the limitation of one floating-point instruction per cycle.
- (2) The decoupled dataflow scheduling becomes extra difficult and inefficient by the variable length of ×86 instructions. Some allocated bits of entries in the instruction window (prefetch buffers) are seldom used and wasted.
- (3) Only single-level speculative execution is supported because of the lack of a result buffer. An unresolved conditional branch instruction at the WB stage will stall the processor and inhibits further speculative execution. Adding a result buffer is not easy because an ×86 instruction can have a memory source operand and destination. Memory-to-register, memory-to-memory, register-to-memory, and immediate-to-memory instructions are permitted. Unlike register identifiers, memory identifiers (addresses) are relatively long. The result buffer will require extra storage and wide comparators for long identifier bits.
- (4) The out-of-order execution is limited to in-order issue with the possibility of out-of-order completion between instructions in the integer and floating-point pipelines. Load bypassing is also not supported.

Summary of Comparisons

Table 3 is a summary of existing commercial superscalar microprocessors (except the Thunder SPARC) discussed in previous sections. The integration, clock speed, SPEC benchmark performance, superscalar features supported, instruction and result shelving type, execution unit composition, instruction decoder size, instruction issue capability, branch prediction type, and notable characteristics of each microprocessor are described. "All six" superscalar features means multi-instruction issue, decoupleddataflow scheduling, out-of-order execution, register renaming, speculative execution, and precise interrupts.

Table 4 is a summary of upcoming commercial superscalar microprocessors in 1995/1996. DEC continues to lead the pack with its new ALPHA 21164 design. The major architectural improvements from its ALPHA 21064 predecessor are quad-issue, additional integer and floating-point execution units (total 2 each), and the inclusion of a secondary cache on chip (see L. Gwennap, "Digital Leads the Pack with 21164," Microprocessor Report, Vol. 8, No. 12, pp. 1 and 6-10, Sep. 12, 1994). The last feature is the first in the history of microprocessors and makes the ALPHA 21164 the densest with 9.3 million transistors. Sun Micro-¹⁵ systems' ULTRASPARC (Trademark of Sun Microsystems Computer Corporation) incorporates nine independent execution units, including dedicated graphics add and multiply units. The ULTRASPARC is the first implementation architecture, which supports conditional move instructions. It also supports MPEG-2 graphics instructions in hardware to boost multimedia application performance. The IBM POWERPC 620 is the latest and currently fastest among other PowerPC models (601, 603, 604, 615). It uses reser-²⁵ vation stations to shelve instructions at six execution units (see L. Gwennap, "620 Fills Out POWERPC Product Line," Microprocessor Report, Vol. 8, No. 14, pp. 12-17, Oct. 24, 1994). IBM put two entries for each execution unit with the exception at load/store unit (3 entries) and branch unit (4 entries). The MIPS Technologies' R10000, also known as T5, uses "decoupled architecture" (another term for

decoupled dataflow scheduling) with a set of three central windows (16-entry queues) for memory, integer, and floating-point instructions (see L. Gwennap, MIPS R10000 Uses Decoupled Architecture," Microprocessor Report, Vol. 5 8, No. 14, pp. 17-22, Oct. 24, 1994). The R10000 uses a register-mapping table (also called rename buffer) to support register renaming. Both integer and floating-point units have 64, 64-bit physical registers that are mapped to 32 logical registers. To handle multi-level speculative execution (up to 4 conditional branches), the R10000 saves the mapping table in shadow registers when encountering a conditional branch. The HP PA-8000 is the first 64-bit PA-RISC architecture implementation (see L. Gwennap, "PA-8000 Combines Complexity and Speed," Microprocessor Report, Vol. 8, No. 15, pp. 1 and 6-9, Nov. 14, 1994). Like its predecessors (PA-7100, PA-7200), the PA-8000 will not have on-chip cache. PA-RISC is the only advanced, general-purpose microprocessor architecture that uses off-chip L1 cache. The AMD K5 is currently the fastest x86 processor, claimed to of the new 64-bit SPARC version 9 instruction-set ²⁰ be at least 30% faster than the Intel PENTIUM at the same clock rate (on integer code-see M. Slater, "AMD's K5 Designed to Outrun PENTIUM," Microprocessor Report, Vol. 8, No. 14, pp. 1 and 6-11, 1994). Despite the CISC ×86 instruction set, the architecture internally runs RISC instructions, called ROPs (RISC operations). To achieve this, ×86 instructions are predecoded as they are fetched from memory to the I-cache. The predecoder adds five bits to each byte, causing an increase of about 50% at the I-cache array. The K5 applies all the superscalar techniques that Johnson believed to be the best, the reservation station for instruction shelving and the reorder buffer for result shelving (see Johnson, 1991).

TABLE 3

	Comparisons of Commercial Superscalar Microprocessors						
	IBM RS/6000	Metaflow THUNDER SPARC	TI SUPERSPARC	DEC ALPHA 21064	HP PA-7100	Intel PENTIUM	
Integra-	Multi-chip	Multi-chip	Single-chip	Single-chip	Multi-chip	Single-chip	
Clock Speed	62.5 MHz	80 MHz (est)	50 MHz	200 MHz	99 MHz	66 MHz	
SPECint92 SPECfp92 Supersca- lar Fea- tures	61.7 133.2 All six	200 (est) 350 (est) All six	68 85 All six, except re- gister renaming	116.5 193.6 Multi-inst. issue, decoupled dataflow sched., out-of-or- der exec.	80.0 150.6 All six, except re- gister re- naming and speculative execution.	64.5 56.9 All six, except re- gister re- naming	
Instruction Shelving	FIFO I-buf- fers (cen- tral, dist. in FXU and FPU)	DCAFs (DRIS) (central, branch, floating- point)	Central win-dow and dist. win- dow (for FP inst.)	Central window (pre-fetch buf-fers)	Central window (prefetch buffers)	Central window (prefetch buffers)	
Result Shelving	Regmap- ping table in FPU	DCAFs	None	None	None	None	
Indepen dent Execution Units	1 branch unit, 1 FX unit, 1 FP unit (MAF)	1 branch unit, 3 ALUs, 1 FP add, 1 FP mult	1 branch unit, 3 ALUs, 1 address add, 1 FP unit (add, mult)	1 branch unit, 1 address unit, 1 integer unit, 1 FP unit (add, mult, div)	1 integer unit (ALU, shift, branch add), 1 FP unit (mult, div/sqrt)	2 ALUs, 1 FP unit add, mult, div)	

TABLE 3-continued

		Comparisons of Commercial Superscalar Microprocessors				
	IBM RS/6000	Metaflow THUNDER SPARC	TI SUPERSPARC	DEC ALPHA 21064	HP PA-7100	Intel PENTIUM
Decode Size Max Issue	4 instruc- tions 1 FX or FP load/store inst., 1 FP arith. inst., 1 branch inst., 1 condi- tion-regi- ster inst.	4 instruc- tions 1 branch inst., 2 integer inst., 1 load/ store inst, 1 FP add/ sub, 1 FP multiply	4 instruc- tions Triple issue with certain restriction s.	2 instruc- tions Dual issue of certain integer/ floating- point operate, branch, and load/store.	2 instruc- tions Dual issue of integer and floa- ing-point instruc- tions.	2 instruc- tions Dual issue of simple instruc- tions.
Branch Prediction	Static (con-stant pre-dicted- not-taken)	Dynamic	Static (al- ways taken)	Static and dynamic	Static (BTFN)	Dynamic
Notes	nFP mult- add in 2 cycles. nThis RS/6000 de- sign is the foundation of follow-on single-chip versions (PowerPC 601, 603, 604, 620) nPrecise interrupts only in synchroniz e mode.	nHas the most com- plete dyna- mic hard- ware sche- duler with full out- of-order issue. nLow clock speed due to complex out-of-or- der issue. nThunder SPARC was a reborn of the unsuccessfu 1 Lightning SPARC.	n Multiple- path fetch- ing into se-quential and target inst. queues helps reduce branch mispredic- tion penalty.	nHybrid of superpipe- lined and supersca- lar. nTrue 64- bit archi- lecture. nSupports cond. move inst. nSupports multiple O/S using PALcode. nImprecise interrupts.	nUses off- chip, pri- mary caches for size and speed flexi- bility nSupports VLIW-like and SIMD- like inst. for path- length reduction.	nSupports the wide- ly-used x86 inst. set. nThe only superscalar processor that is not a register- to-regi- ster, 3-ad- dress ma- chine. n Inefficient inst. window due to vari- able-length x86 inst.

TABLE 4

	Comparisons of Upcoming Comercial Superscalar Microprocessors							
	DEC ALPHA 21164	Sun Micro- systems ULTRASPAR	IBM POWERPC 620	MIPS Technolo- gies R10000 (TS)	HP PA-8000	AMD K5		
Integra-	Single-chip	Single-chip	Single-chip	Single-chip	Mult-chip	Single-chip		
Clock Speed	300 MHz	167 MHz	133 MHz	200 MHz	n/a	n/a		
SPECint92 SPECfp92 Supersca- lar Features	330 500 Multi-inst . issue, de-coupled da-taflow sched., out-of-or- dar swap	275 (est) 305 (est) All six, except re- gister renaming	225 (est) 300 (est) All six	300 (est) 600 (est) All six	n/a n/a All six	n/a n/a All six		
Instruc- tion Shelving	Central window pre-fetch buffers)	Central window and dist. win- dow (for FP inst.)	Reservatio nStations	Central/ dist. win- dow (mem- ory, inte- ger, FP queues)	Central window (56 entries)	Reserva- tion Sta- tions		
Result Shelving	None	None	Rename buf-fers (reg. map	Register- mapping table	Rename buf-fer	Reorder buffer		

TABLE 4-continued

	Comparisons of Upcoming Comercial Superscalar Microprocessors					
	DEC ALPHA 21164	Sun Micro- systems ULTRASPAR	IBM POWERPC 620	MIPS Technolo- gies R10000 (TS)	HP PA-8000	AMD K5
Indepen- dent Execution Units	1 branch unit, 1 address unit (load/ store), 2 integer units, 2 FP add/ div, 1 FP mult	1 branch unit, 1 address unit, 2 ALUs, 1 FP add, 1 FP mult, 1 FP div/ sqrt, 1 graphics add, 1 graphics	table) 1 branch unit, 1 address unit, 2 ALUs, 1 complex integer unit, 1 FP unit	1 branch unit, 1 address unit, 2 ALUs, 1 FP add, 1 FP mult/ div/sqrt	1 branch unit, 2 address units, 2 ALUs, 2 FPUs (MAC /div/sqrt)	1 branch unit, 2 ALUs, 2 load/ store, 1 FPU
Decode Size Max Issue	4 instruc- tions Quad issue of 2 integer and 2 floating- point in- structions.	4 instruc- tions Quad issue	4 instruc- tions 6 (w/ dist. windows, each exec. unit can issue an inst.	4 instruc- tions 6 instruc- tions (1 instructio nfrom each execution unit)	4 instruc- tion 4 instruc- tions (2 to address unit, 2 to ALUs/ FPUs)	4 instruc- tions 6 ROP in- structions (1 inst from each execution unit)

Branch Predic- tion	Static and dynamic	Dynamic	Dynamic	Dynamic	Dynamic	Dynamic
Notes	 Towering manufactur -ing cost (\$430) leads to an expensive initial price of \$2,937 (9/12/94) Most dense microprocessor to date at 9.3M transistor s 	 Include a special purpose graphics unit that supports MPEG-2 in- structions First 64-bit V9 instructio n-set archi- tecture im- plementa- tion 	• Based on SPECint92, the 620 design does not improve much from 604, par- ticularly if the 604 reaches 133 MHz also. Key advan-tage of 620 is its higher mem-ory band- width.	• Has a "resume" cache to hold se- quential alternate- path in- structions in case of branch mis- prediction • Handles 4-level speculativ e execution	• First 64-bit PA- RISC archi- tecture im- plementa- tion • Like its predeces- sors, no on-chip cache, off- setting penalty by having large off- chip L1 cache	 Claimed to deliver at least 30% faster than PEN- TIUM at the same clock rate (on integer code). On the average 16-bit x86 code produces 1.9 ROPs/ inst., 32- bit x86 code pro- duces 1.3 POPe/inst.

Motivation for the Invention

The primary objective of the invention is to find a better design alternative to the reorder buffer (considered the most complete, known result shelving technique) and the reservation station (considered the best known instruction shelv- 60 ing technique to give maximum machine parallelism) techniques. Before pointing out to the drawbacks in the two best, existing techniques in result and instruction shelving, let us understand fully how the two operate in a superscalar processor. FIG. 18 illustrates how the RB and the RS 65 These instructions are decoded in parallel and dispatched to techniques can be combined to fully support all six superscalar features.

Assume the processor has eight execution units: a branch unit, two fixed-point ALUs, a floating-point ALU, a floatingpoint multiply/divide/square-root unit, two load/store units, and a fixed/floating-point move unit. The processor begins by fetching at least one instruction or multiple instructions $(N_d \text{ instructions in this case, which is the decoder size})$ from the I-cache. It is to be understood that one or more memories or other storage units may be employed instead of the I-cache for performing the same function as the I-cache. their respective execution unit's reservation station. For each decoded instruction, an entry is allocated at the RB to

55

shelve its result. To read its operand(s), each operand's register number is presented to the RB and register file, in which three situations can occur. First, if there is a matched entry in the RB and the register operand value has been calculated, then the operand value is routed/copied to the instruction's reservation station. Second, if there is a match entry in the RB but the value has not finished calculation, then the operand tag is copied instead to the reservation station. Third, if there is no match entry in the RB then the value from RF is the most recent one and copied to the 10 reservation station.

During the issue/execute stage, one ready instruction (i.e., all of its operand values are available) is issued and executed immediately, even if it is out of order, and independent of other execution units' issues. If there are multiple ready 15 instructions, one is chosen based on the arbitration protocol (aging, random, pseudo-random, small index, etc.).

During the writeback stage, the execution result is written to the RB (not RF) and also forwarded to any reservation 20 station waiting for this result value. In every cycle, each valid reservation station with unavailable operand(s) compares its operand tag(s) with all result tags to determine when to grab certain result value(s). Note that if each execution unit's output port does not have a dedicated result 25 bus (N_{res}<fixed-point output ports, or N_{resfp}<floating-point output ports), then arbitration logic must be provided to resolve who can use the shared result buses at a given time.

During the retire stage, the bottom entry at the RB FIFO is checked to see if the value can retire to the RF or not. An instruction's RB entry can only retire if: (a) the instruction completed execution without an exception error, (b) all preceding instructions also completed without exception, and (c) the instruction is not speculative (beyond an unresolved conditional branch). To keep up with the input rate at the fetch/decode stage, multiple entries (N_{ret}) must be retired, in order, from RB to RF.

Drawbacks in the Reorder Buffer Technique

The major implementation drawback in the RB technique 40 is the slow and expensive prioritized associative lookup. During operand read in the decode stage, a register number is presented to the RB to find the most recent entry that matches the register number. However, the register number is not a unique associative key. Often there are multiple 45 entries in the RB with the same register number, due to register re-uses (multiple updates to the same register of different instructions). In such a case, the associative lookup must be prioritized so that the most recent entry (of the most recent assignment) is selected from the RB and also from 50 preceding instructions in the decode group that are currently in RB allocation ports (about to be written). If the RB is implemented as a circularly addressed register array, then the prioritizing circuit is complicated by the constant change of the head and tail pointer positions with the possibility of a wrap around. The extra delay incurred by the prioritizing logic will slow down operand reads.

The prioritizing circuit complexity can be reduced by constructing the RB as a true FIFO array. However, the FIFO array implementation has a penalty in supporting multi-entry 60 retire. A FIFO typically shifts out one entry at time. With multiple-instruction fetch and decode, multiple RB entries in the FIFO must be deallocated/freed to keep up with the allocation rate and reduce decode stalls due to RB being full. To shift multiple entries at a time, two things can be done. 65 First is to clock the shift-register array using a frequency that is N_{ret} times the processor clock frequency. This is imprac-

tical because reads from and writes to the RB are very difficult and incur more propagation delay due to changing FIFO entry positions in a processor cycle. At the end the processor cycle time will just get longer. Second is to implement the FIFO as a barrel-shift register array with a variable number of shifts from 1 to N_{ret} , as shown in FIG. 19. Now each entry at position i must be able to mux in from any entry above it at positions i+1, i+2, ..., i+N_{ret} or the last FIFO array index. This barrel-shift register array implementation requires a substantial number of multiplexers.

Besides the prioritized associative lookup problem, the RB has two other drawbacks, the need to allocate dummy branch entries (entries without a result/register value) and the substantial use of shared global buses. To support multi-level speculative execution, a typical RB allocates a dummy RB entry for every decoded conditional branch instruction. This increases the usage of RB entries unnecessarily. The RB technique also requires that operand values to be distributed to all participating functional units, resulting in a large amount of global (chip-wide) wire traces that are heavily loaded. The result buses that carry result values to be written back to the RB also require shared-global wires that are heavily loaded with multiple driving sources, which are prone to transmission line noise problem. With the continuing shrinkage of transistor sizes, global wires are becoming more dominant in terms of silicon area consumption, propagation delay, and signal noise.

The complexity of RB circuit implementation is a proven fact. AMD found out in their superscalar 29K processor (see 30 B. Case, "AMD Unveils First Superscalar 29K Core," Microprocessor Report, Vol. 8, No. 14, pp. 23-26, 1994), an "intellectual predecessor" to the K5, that a mere 10-entry RB plus its operand and result buses consumed a significant chip area, about the same as a 4-Kbyte cache, with the chip 35 fabricated in an advanced 0.5-micron CMOS technology!

Drawbacks in the Reservation Station Technique

There are two major drawbacks in the RS concept; (1) operand value copying and (2) result forwarding. Both result in tremendous amounts of shared (heavy-loading), global (chip-wide) wires. With increasingly smaller transistor sizes, the dominant factors in determining silicon area and propagation delays is not the transistor, but metal wire, especially the ones that run across or all over the chip. Therefore, we should view the global wire as a precious item in designing a circuit, in particular if it is a shared line with many fanouts or loadings.

The wide operand buses required for operand value copying waste a lot of precious data bandwidth in the processor chip. Most of the time only a few of the total operand buses $(2*2*N_d*64 \text{ global wires for a 64-bit})$ architecture) are used. The DLX simulation results show a low bandwidth usage of less than 7% for the operand buses! The waste is due to several factors:

- Many decoded instructions have only one register operand (arithmetic instructions with immediate value operand, loads, conditional branches, floating-point convert instructions), or worse, no register operands at all (jumps, traps).
- When there is a branch-taken instruction (either predicted-taken conditional branch or unconditional jump) in the decode group, the subsequent instructions are invalidated, resulting in zero operand bus usage.
- When a register operand value is not available at the RB (still being calculated), the operand bus only carries a small number of operand tag bits (e.g., a 6-bit tag in a 64-bit data bus uses only 9% of that bus bandwidth).

When the decode stage stalls, zero operand buses are used.

Besides the expensive global wire cost, result forwarding also requires complex circuitry. Each reservation station must have comparators to associatively match any tag in 5 result buses, and wide multiplexers (64-bit, [2*Na+Nres]:1 muxes) to receive operand values not only from operand buses, but also result buses.

Result value forwarding also requires a set of additional comparators (exactly $2*\bar{N}_d*N_{res}+\bar{2}*N_d*N_{resf}$ comparators) to avoid potential pipeline hazards in the forwarding logic (see Weiss and Smith, and Johnson, 1991). The hazard arises when a decoded instruction's operand tag matches with the result tag of a result value that will be forwarded at the same time the decoded instruction is dispatched to a reservation station. Had this hazard not been detected, the instruction 15 will miss the forwarded result and stall in the RS for a long time until the tag has wrapped around, causing an incorrect operation. FIG. 20 illustrates what happens if this resultforwarding hazard is not detected. The example shows that instruction I_r is decoded at cycle t, executed on the first half 20 of cycle t+1, and its result is forwarded or appears at a result bus on the second half of cycle t+1. In the meantime, instruction I,, which has a data dependency with I, on result R3, is decoded on the first half of cycle t+1 and dispatched to a RS with R3's tag on the second half of cycle t+1, not 25 knowing that the R3 value is being forwarded at this time. At cycle t+2, I, tries to execute but its R3 operands are unavailable, hopelessly waiting for the missed R3 value to appear at one of the result buses. To overcome this resultforwarding hazard problem, we can simply stall the decode 30 stage with some performance penalty, or provide a set of wide-bus multiplexers to allow operand value bypassing from the result buses. Under normal condition, a RS receives operand values from the register unit. When a resultforwarding hazard is detected, then the RS can receive the correct operand value directly from the corresponding result 35 bus.

Proposed Solution to Drawbacks in the Reorder Buffer and Reservation Station Techniques

The invention involves a design alternative that will solve 40 the drawbacks in the reorder buffer and reservation station techniques.

The invention presents new or improved result shelving and instruction shelving techniques that maintain the good characteristics of reservation station and reorder buffer 45 designs, but eliminate or reduce their major implementation disadvantages. By maintaining/improving the good characteristics in the new design, we can reasonably target a performance that is on par or better on a cycle-by-cycle basis (cycle count). By eliminating or reducing the implementa- 50 a more flexible way of assigning unique result tags, which tion drawbacks we ensure that the hardware cost decreases (smaller circuit/die area) and the actual speed or performance of the processor increases (faster circuit leading to reduced cycle time or more operations done per cycle). However, care must taken not to propose a hardware sim- 55 plification which could result in some penalty on the cycle count. Intelligent choices must be made to minimize the cycle count penalty such that it can be easily recovered through faster and more efficient circuit implementation.

The following sections describe the invention, the Modi-⁶⁰ fied Reorder Buffer (MRB) and the Distributed Instruction Queue (DIQ. An example of how the MRB and DIQ are used together in a superscalar processor is given below.

Modified Reorder Buffer (MRB)

Before the MRB technique is presented, the thinking process that led to the proposed MRB solution is described.

Then the MRB circuit structure is explained. To support multi-level speculative execution without the need to enter dummy branch entries in the RB, a new small circuit, the Flush Address Queue (FAQ) is added and described.

Rationale for the MRB

The reorder buffer was originally proposed by Smith and Plezkun to handle precise interrupts in pipelined, scalar processors with out-of-order completion. Conceptually, the RB is a content-addressable, circular buffer with head and tail pointers. To use Smith and Plezkun's RB "as is" in superscalar processors will produce unreasonable hardware cost and seriously limit performance because of the "multiple-bypass check" (prioritized associative lookup) problem (see Smith and Pleszkun), the limited access ports (only 1 allocate port, 2 read ports, 1 write port, and 1 retire port) that impedes machine parallelism, and the inability to support speculative execution. This RB paper design was modified by Johnson to make it more suitable for superscalar processors (see Johnson, 1991), and later realized in the superscalar 29K and AMD K5 processors (see Slater, Case). The modifications include: (a) providing multiple ports for allocate (N_d ports), read (2* N_d ports), write (N_{res} ports) and retire (N_{ret} ports), (b) allocating a dummy entry for every conditional branch instruction to support speculative execution, and (c) eliminating the unnecessary "result shift register" to reserve a result bus, since the reservation stations already contain an arbitration mechanism for result buses.

Although Johnson's RB is probably considered the most complete result shelving technique, it still has one major implementation drawback, the prioritized associative lookup This problem originates from a flaw in the RB concept, the use of a non-unique associative key (register number) during RB's associative read. This non-unique associative key forces the RB to prioritize the associative lookup, which results in a slow and complex hardware. Therefore it is logical that we need to replace this register number with a unique associative key.

Ideally we want to use the same result register tag, which is used during the result write operation, as the unique associative key for read operation. This tag is written to the RB entry during allocation. Smith and Plezkun use the RB identifier or array index as the result tag. But this tag is not unique with the presence of a second RB (e.g., for a floating-point register file). Moreover, the tag will keep changing as the FIFO queue is advanced during multi-entry retire. Tracking many different register tags plus conditional branch tags can be a nightmare. Weiss and Smith suggested was originally proposed to be used in reservation stations (see above section entitled Reservation Stations). When an instruction is decoded, a new tag or identifier (inst_ID) is assigned from a "tag pool" that consists of some finite set of tags. Each destination register is then tagged with the inst_ID of the producer instruction. When the instruction completes, the inst_ID is returned to the pool for reuse. This tag pool, called the Instruction ID Unit (IIU), can be implemented as a circular instruction array (IA). The inst ID is composed of (color_bit, IA_index-these are discussed in more detail below), the current "color" bit appended with its IA index (entry address), the same scheme used in the DRIS technique (see above section entitled DRIS). The color bit is used to distinguish the age or order 65 of instructions when the valid entry area wraps around.

Now the question is, how do we get the source operand tag to be used as the unique associative key when reading the

35

40

65

RB, without reading the RB first? Remember that each decoded instruction with register operands needs operand tags (or operand values in the case of a reservation station technique) before being dispatched to an instruction window. To accommodate these operand tags at the decode stage, a Register Tag Unit (RTU) is added in the fetch and decode unit. Each decoded instruction presents its source register numbers to the RTU to get the corresponding register tags. The RTU can be viewed as a small register file that maintains the most recent tag (not value) of every register. The most recent tag of a register is the inst_ID of a producer instruction that updates the register last. When an instruction is assigned an inst_ID by the IIU, the destination register entry in the RTU is written with the inst_ID.

With minimal hardware support in the fetch and decode 15 unit (IIU and RTU), we can now construct the modified reorder buffer (MRB) that uses a unique associative key for read, write, and retire operations. With this unique associative key we can implement the MRB as a simpler, circularly addressed register array instead of the more expensive true 20 FIFO array, because we do not have to prioritize the associative read/lookup anymore! Moreover, with circular array implementation, multiple-entry retire is simply done by moving the head pointer several locations up. Therefore, we have eliminated one implementation drawback in a typical 25 RB.

MRB Structure

FIG. 21 shows the MRB 100 structure with N_{xrb} entry cells 195, N_d allocate ports 110, N_{rd} read ports 130, $N_{wr=30}$ write ports 120, and N_{ret} retire ports 140. An MRB entry consists of six fields 185: not empty flag 186 (1 bit), valid flag 187(1 bit), exception error flag 188 (1 bit), register number 189 ($\log_2 N_{xrf}$ bits, N_{xrf} is fixed-point RF size), register tag 190 (N_{tag} bits, N_{tag} =1+ $\log_2 N_{ia}$, N_{ia} =tag pool/ array size), and register value 191 (N_{data} bits) Each MRB entry cell 195 consists of D flip-flops (DFFs) or RS flip-flops (RSFFs)105 to hold the MRB entry fields 185 and logic circuits for allocate 145, read 125, write 115, and retire operations 135.

During the decode stage, an entry is allocated to MRB 100 for each register assignment (R_alloc_en (i)=1). The setting of a new MRB 100 entry is: n_empty=1, valid=0, excp=0, reg_num=destination register number, and reg_ tag=inst_ID assigned to the instruction. There can be up to 45 N_d entries allocated in a cycle. To prevent overlapping/ overflowing during allocation, the n_empty flag must be 0 (indicating the MRB cell 195 is empty) before an allocation is made. overlapping must be prevented because of the circular buffer nature. If MRB 100 overflows, a fetch and 50 decode unit 520 is notified immediately so that the fetch and decode stages are stalled until there is room in MRB 100.

During the issue stage, operand tags are presented to MRB 100 from read ports 130 to read their register values from register value field 191. If a matched entry is found 55 (reg_tag=operand tag), the register value has been calculated (valid=1), and the producer instruction completes without exception (excp=0), then the reg_value is passed to an execution unit (505, 506, 560, 570, 580, 590, 595, or 596—see FIG. 31). If the MRB entry at MRB 100 is found 60 but the register value is not valid (valid=0 or excp=1), then the issue is put on hold or stalled. If no entry is found at MRB 100, then the register value is in RF 515R or 515F. This means RF 515R or 515F contains the most up-to-date value of the operand.

During the writeback stage, as the instruction completes its execution, the result value is written to the reg_value

field 191 of the previously allocated MRB entry cell 195. This MRB entry cell 195 can be found because every result is accompanied by its unique producer inst_ID (reg_tag= result tag). If the instruction completes with an exception, the excp bit is set. However, the processor (may be processor 500 in FIG. 31) is not interrupted immediately but delayed until the retire stage. This ensures precise interrupts and prevents the processor from servicing false exceptions caused by: (a) a speculative instruction in a mispredicted path, or (b) an out-of-order instruction that completed early 10 with an exception, but is then followed by a second exception of an older instruction $(2^{nd} \text{ exception's inst_ID} < 1^{st})$ exception's inst_ID). This MRB entry cell 195 with a false exception will eventually be flushed before reaching the bottom of MRB 100, due to branch misprediction's flush or another exception's flush.

During the retire stage, N_{ret} entries at the bottom of MRB 100 are checked to see if some or all can retire to the corresponding RF 515R or 515F. When the retire operation occurs, the reg_value 191 field values are retired from MRB entry cells 195 to the corresponding RF 515R or 515F (see FIG. 31). The conditions for the field 191 value in an MRB entry cell 195 to retire are:

- (1) the producer instruction has completed without an exception (excp=0) and the result value has been written (valid=1),
- (2) instructions preceding the producer instruction also completed without an exception (reg_tag<in_order_ point), and
- (3) the producer instruction is non-speculative (reg_ tag
branch_point).

The in order point (generated by fetch and decode unit 520) is the inst_ID of the "oldest" instruction that has not completed or completed with an exception. Thus, any instruction with inst_ID<in_order_point, has completed successfully without an exception. The branch_point (generated by the branch unit—not shown) is the inst_ID of the "oldest" unresolved conditional branch. Therefore, all instructions prior to the conditional branch (inst_ ID>branch_point) are non-speculative.

The head pointer logic unit or block (head_RB)170 generates retire locations and adjusts the head pointer 180 positions for the next cycle. Since all N_{ret} entries at the bottom of MRB 100 are simultaneously evaluated to retire, head_RB(i)=(head_RB(0)+i) mod N_{xrb} is generated for each i ϵ [0,N_d-1]. Each valid retirement (valid_ret(i)=1) moves the head pointers 180 "up" by one position. So, the next cycle head pointer 181 is,

(next)
$$head_RB = \left(head_RB + \sum_{k=0}^{N_{ret}-1} valid_ret(k)\right) \mod N_{xrb}.$$
 (2)

The tail pointer logic unit or block (tail_RB)150 determines the correct tail pointer 160 position of every allocation port 110:

$$tail_{\mathbf{RB}}(i) = \left(tail_{\mathbf{RB}}(0) + \sum_{k=0}^{i-1} R_{\underline{alloc}_{\underline{en}}(k)}\right) \mod N_{xrb},$$
(3)

where i $\in [1, N_d - 1]$. The increment from the base tail pointer 160 is determined based on the number of allocations (how many R_alloc_en(k)=1 for k<i) in the previous allocate ports 110. The next cycle's tail_RB(0) is set to (tail_RB $(N_d-1)+R_alloc_en(N_d-1) \mod N_{xrb}$, provided no branch misprediction is detected (mispred_flag=0). If mispred_ flag=1, then the next cycle tail_RB(0) is set to the RB_flush_tail from the RB Flush Address Queue (RBFAQ) 165.

FIG. 22 shows a circuit implementation of MRB 100, 5 with the cell circuitry shown in FIG. 23 and the support logic (tail pointer logic 150, head pointer logic 170, and flags logic 155) shown in FIG. 24. As shown in FIG. 22, each MRB cell 195 comprises RSFFs or DFFs 105, write logic 115, read logic 125, retire logic 135, allocate logic 145, and mispre- 10 diction logic 175. In FIG. 23 Retire logic 135 has comparison unit 152 which may be a comparator or other device capable of comparing inputs as is well-known to those of ordinary skill in the art for checking to see if it is safe to retire and misprediction logic 175 has comparison unit 151 15 which may also be a comparator or other device capable of comparing inputs as is well-known to those of ordinary skill in the art for flushing misspeculative MRB 100 entries. Throughout all schematics, a thick line indicates a multi-bit bus (multiple wires), while a thin line indicates a single-bit 20 wire. For notational consistency, the following signal name conventions are made. A signal named "X(i)" indicates that the signal is the ith bus in a group of buses named X. A signal named "Y[k]" indicates that the signal is the kth bit of bus Y. A signal named " $Z(i)_z[k]$ " means that the signal is the 25 k^{th} bit of bus Z_z, which is the i^{th} bus in the group bus Z. A signal name "X[k]" can also denote a register output name (at position k) if it is part of a register array.

The following section describes how multi-level speculative execution can be supported efficiently with the addi- 30 tion of RBFAQ 165.

Flush Address Queue (FAQ)

To support speculative execution, a mechanism to recover and restart must exist so that correct results can still be 35 produced in the event of branch misprediction. Recovery cancels the effect of instructions under false predictions, and restart re-establishes the correct instruction sequence. When a branch misprediction is detected, speculative RB entries (that turn out to be mis-speculated) must be flushed to nullify the effect of wrong updates. The main difficulty is to quickly locate the starting speculative entry (reset point) for

starting flush tail/address of each unresolved conditional branch. RBFAQ 165 can be implemented as a true FIFO, or better, a circular array. Each RBFAQ 165 entry in entry cells 295 consists of only log₂N_{xrb}-bits of MRB tail pointer 160 (flush_addr). When a conditional branch instruction is decoded (e.g., at decode slot i), an entry is allocated at the tail of RBFAQ 165, with the entry's flush_addr set to the current tail_RB position, tail_RB(i). Up to N_d entries can be allocated per cycle. When the "oldest" conditional branch instruction completes at the branch unit, the misprediction flag (mispred flag) is checked. If mispred flag=0 (no misprediction), then the corresponding entry in RBFAQ 165, which is obviously at the bottom entry cell 295, is simply discarded/popped. If mispred flag=1 (misprediction), all speculative entries in MRB 100 are flushed by simply setting tail_RB(0) to the old tail position when the conditional branch was decoded (RB_flush_tail=entry's flush_addr). RBFAQ 165 is then reset (completely flushed) because all subsequent entries belong to branches in the mispredicted path. The number of entries in RBFAQ 165 (N_{faq}) defines the maximum levels of speculative execution (total unresolved conditional branches) supported.

The FAQ tail pointer determined by tail pointer logic 250 (FIG. 25) for an allocation at decode slot i (FAQ_alloc_ en(i)=1 if there is a valid conditional branch instruction) is calculated as follows:

$$tail_FAQ(i) = \left(tail_FAQ(0) + \sum_{k=0}^{i-1} FAQ_alloc_en(k)\right) \mod N_{faq},$$
(4)

where i ϵ [1,N_d-1]. The increment from the b base tail pointer is determined based on the number of allocations in the previous ports. The next cycle tail_FAQ(0) is set to $(tail_FAQ(N_d-1)+FAQ_alloc_en (N_d-1)) \mod N_{faq}$. However, if a branch misprediction is detected, the tail_ FAQ(0) is set back to zero. There is only one head pointer 40 determined by head pointer logic 270 which is updated as follows:

(5) if branch completes w/mispred_flag = 1 $head_FAQ = \begin{cases} (head_FAQ + 1) \mod N_{faq}, \text{ if branch completes w} / mispred_flag = 0\\ head_FAQ, & \text{otherwise (unchanged)} \end{cases}$

60

65

flushing, so that recovery and restart can be done with minimum delay. Johnson suggests allocating a "dummy' branch entry in the RB for every conditional branch instruction decoded (see Johnson, 1991). The branch entry is dummy because there will be no register update or result value produced. When the conditional branch is detected to 55 be mispredicted, processor searches the corresponding branch entry in the RB. This entry becomes the reset point; all entries subsequent to it are flushed by setting the tail pointer to the branch entry number. Or alternatively, we can let the "branch" entry reach the bottom (head) of the RB. However, this adds delay to the recovery and restart from a branch misprediction. The main disadvantage with the dummy-branch-entry method is the increase in RB usage. There can easily be multiple unresolved branch entries in the RB that increases the number of decode stalls due to RB full/overflow, or forces more entries in the RB.

As a solution to this problem, a small RB Flush Address Queue (RBFAQ)165 is added to MRB 100 to hold the

An example of the RB flush address queue (RBFAQ)165 circuit implementation, comprising DFFs 205 or other equivalent devices, is shown in FIG. 25, with supporting logic in FIG. 26. Supporting logic for RBFAQ 165 also includes flags logic 255 for d determining the state of RBFAQ 165.

Distributed Instruction Queue (DIQ)

The DIQ is a new distributed instruction shelving technique that offers a solution to drawbacks in the reservation station (RS) technique. First, a rationale that led to the DIQ solution is described. Then, the in-order DIQ 300 circuit structure is explained. Finally, an enhancement to the DIO, the out-of-order DIQ 400 that will allow full out-of-order issue is described.

Rationale for the DIQ

First, based on comparisons of the DRIS central window and RS distributed window, it was decided early that the

proposed instruction shelving should be a distributed window technique. A distributed window implementation has advantages over a central window implementation because of its

- maximum machine parallelism (all execution units can 5 issue an instruction simultaneously),
- simpler, independent issue logic (issue only one instruction locally from a relatively small group of instructions, independent of other windows' issues),
- many-to-many routing paths),
- simpler deallocation logic (free only one window entry at a time), and
- more efficient use of the window (deallocate issued entry immediatelv).

To compensate for the implementation complexity of a central window, most commercial processors with a central window method have to split the large central window into smaller sub-central windows (see Table 4 and 5). This half-way solution still has most of the central window 20 the simulation results. drawbacks.

The reservation station (RS) technique, currently considered the best technique, was originally introduced by Tomasulo in 1967 in the floating-point section of the IBM 360/91 (see Tomasulo). The main objective was to permit simulta- 25 neous execution of independent instructions while preserving the precedence (data dependency) constraints in the instruction stream. Tomasulo's RS technique was essentially ahead of its time. It actually accomplishes several superscalar objectives; multi-instruction issue, decoupled dataflow 30 scheduling, out-of-order execution, and register renaming (eliminating anti and output dependencies). Anti dependencies (write-after-read hazards) are avoided by tagging registers, copying operands to reservation stations, and dependencies (write-after-write hazards) are avoided by comparing tags at the FLR (floating-point register unit in the IBM 360/91) on every register write, to ensure that only the most recent instruction changes the register. However, Tomasulo's algorithm lacks a mechanism to handle specu- 40 lative execution. Only the most recent updates of registers are maintained, regardless of whether they are speculative updates or not. To support multi-level speculative execution, the register file can be accompanied by a reorder buffer (RB) Slater), or multiple copies of register-mapping tables (RMT) as seen in IBM POWERPCs and MIPS R10000 (see Gwennap, Oct. 24, 1994).

As we can see above, there is an overlap of task since both the RS and RB/RMT support register renaming. And we 50 know that RS's efforts to support register renaming by operand value copying and result value forwarding actually penalize the RS implementation, due to excessive sharedglobal wires, comparators, and multiplexers. Therefore, if we already have RB or RMT or MRB result shelving, it 55 seems logical to eliminate the expensive operand value copying and result value forwarding concept. The key concept is that no register values are stored in the DIQ, only their register tags. This eliminates large amount of global buses and wide-bus multiplexers or tristate buffers for data 60 routing. Operand values are read directly during the issue/ execute stage from the reorder buffer or register file, when they are available. Unlike the RS technique which reads operand values during decode stage, the DIQ technique does not have the result forwarding hazard problem. Therefore, 65 we save a substantial number of comparators and wide-bus multiplexers (no operand value bypassing).

Ready instructions can be issued from the DIO in different fashions. First, the simplest way, is to allow only in-order issue within a functional unit. Out-of-order executions can still be achieved amongst execution units (partial out-oforder issue with out-of-order completion). Second is to allow some form of limited out-of-order issue such as checking the next entry if the bottom entry has waited for some time (could be as short as one cycle), or to read operands of the bottom two entries simultaneously. Third is simpler instruction issue routing (many-to-one instead of 10 to allow full out-of-order issue by keeping track of the ready operand tags. The selection of the issue logic protocol can vary from one functional unit to another. For example, the branch unit's DIQ which stores all pending conditional branch instructions may be better off to use in-order issue 15 logic. But the fixed-point ALU's DIQ could benefit from the out-of-order issue logic. Extensive simulations must be performed across a good set of benchmark programs that represent the intended applications. The quality of the compiler and its optimization options will significantly influence

In general, the in-order issue (simplest way) may be good enough due to several reasons. First, the performance loss could be recovered through faster circuit implementation by the simplicity of in-order versus out-of-order issue logic. Second, this problem could be easily solved in software by reordering instructions using list scheduling, control dependence analysis, software pipelining, etc. If instructions were ordered or entered properly in the distributed instruction window, issuing instructions out of order within an execution unit does not increase performance significantly. Third, even without software scheduling, Johnson's extensive trace-driven simulations show that the performance loss of in-order issue versus out-of-order issue reservation stations is negligible in most cases. Across the entire range of forwarding results directly to reservation stations. Output 35 reservation-station sizes, the biggest loss in average performance is 0.6% for a two-instruction decoder ($N_d=2$) and 2% for a four-instruction decoder (N_d =4) (see Johnson, 1991). (The real-application benchmark programs used were ccom-optimizing C compler), irsim-delay simulator for VLSI layouts, troff-text formatter for typesetting device, and yacc-compiles a context-free grammar into LR(1) parser tables. These programs were compiled into MIPS instructions.) The analogy is that, when a dependency stalls instruction issue at a particular execution unit, it is more as seen in the AMD superscalar 29K (see Case) and K5 (see 45 important to issue instructions at other execution units (which will free the stall) than to issue a newer instruction at the stalled execution unit (see Johnson, 1991).

> With DIQs, we can efficiently construct distributed instruction windows that gives maximum machine parallelism, but eliminates completely both implementation drawbacks in the RS concept, operand value copying and result value forwarding. The DIQ technique reduces the number of shared-global wires, comparators, and multiplexers significantly.

In-Order Issue DIQ Structure

FIG. 27 shows in-order issue DIQ 300 structure with N_{dig} entry cells 395 and N_d allocate ports 310, implemented as a circularly addressed register array. It has multiple allocate ports 310 and a single issue port 340. The DIQ 300 entry fields 385 vary from one execution unit to another with the first two fields, 386 and 387 (inst_ID and opcode) always present. The example shown is of a floating-point ALU which consists of: instruction tag 386 (Ntag bits), opcode 387 $(N_{opc} \text{ bits})$, source 1 register number 388 $(\log_2 N_{frf} \text{ bits}, N_{frf})$ is floating-point RF size), source 1 register tag 389 (N_{tag} bits), source 2 register number 390 (log₂N_{frf} bits), and

15

20

35

40

45

60

57

source 2 register tag 391 (N_{tag} bits). Note that in contrast to MRB 100 which only retires field values in field 191 of appropriate cells 195 in the retire operation, DIQ 300 issues all field values of fields 386, 387, 388, 389, 390, and 391 of a single cell **395** in order during the issue operation to instruction issue register 516 (see FIG. 33). Each DIQ entry cell 395 consists of D flip-flops (DFFs) or other storage units 305 (FIG. 28) to hold these DIQ entry fields and logic for the allocate operation as determined by allocate logic in DIQ cell 395.

At the end of the decode stage, the fetch and decode unit (may be 520—see FIG. 31) dispatches the decoded instructions to all execution units (may be 560, 570, 580, 590, 595, 596, 505, or 506 in FIG. 31). Each execution unit monitors the FU_num field (execution/functional unit number) in all dispatch buses (may be 550 in FIG. 31). If the FU_num matches with the execution unit's assigned FU_num, it allocates a DIQ entry for that decoded instruction. The allocate logic 345 generates the control signal to enable the DIQ 300 allocation,

$$DIQ_alloc_en(i) = \begin{cases} 1, \text{ if } FU_num(i) = \text{ execution unit's } FU_num & (6, 0, 0) \\ 0, \text{ otherwise} \end{cases}$$

where i $\in [0, N_d - 1]$. Up to N_d entries can be simultaneously 25 allocated to DIQ (300). If DIQ 300 overflows, the fetch and decode unit is notified immediately so that the fetch and decode stage are stalled until there is room in DIQ 300. The tail pointer logic unit or block (tail_DIQ)350 determines the correct tail pointer **360** position of every allocation port **310**: 30

$$tail_DIQ(i) = \left(tail_DIQ(0) + \sum_{k=0}^{i-1} DIQ_alloc_en(k)\right) \mod N_{diq},$$
(7)

where i $\in [1, N_d - 1]$. The increment from the base tail pointer **360** is determined based on the number of allocations in the previous ports. The next cycle tail_DIQ(0)361 is set to $(tail_DIQ(N_d-1)+DIQ_alloc_en(N_d-1)) \mod N_{dia}$, provided no branch misprediction is detected (mispred_flag= 0). If mispred_flag=1, then the next cycle tail_DIQ(0)361 is set to the DIQ_flush_tail from the DIQ Flush Address Queue (DIQFAQ)365. This essentially flushes instructions in the mispredicted branch path (if any). DIQFAQ 365 is identical to the one used MRB 100, providing multi-level speculative execution. Instructions are issued in order from the bottom of DIQ 300, pointed by issue pointer 380 as determined by head (issue) pointer logic 370. Issue pointer **380** is equivalent to a head pointer (issue_DIQ=head_DIQ) and therefore may also be designated as head pointer **380**. If there is an instruction (DIQ_empty=0), its register operands are read from the result shelf or directly from the register file. If both reads are successful (valid_read(L)=1 and valid read(R)=1) then the instruction is issued for execution, and DIQ 300 is popped. The DIQ head pointer ⁵⁵ 380 is then advanced by one position,



$$\begin{cases} (head_DIQ+1) \bmod N_{diq}, & \text{if } (!DIQ_empty \& valid_read(L) \& valid_read(R) \\ head_DIQ, & \text{otherwise (wait, no issue)} \end{cases}$$

FIG. 28 shows an example of a DIQ 300 circuit imple- 65 mentation according to the preferred embodiment, with the supporting logic (allocate logic 345, flags logic 355, tail

pointer logic 350, and head pointer logic 370) shown in FIG. 29. The DIQFAQ 365 circuit is identical to the RBFAQ 165, except the role of tail_RB is replaced by tail_DIQ. Note that in FIG. 29 tail pointer logic 350 comprises comparison unit 351 which may be a comparator or equivalents thereof as would be understood by those of ordinary skill in the art. Comparison unit 351 is used for flushing misspeculative DIQ entries.

Out-of-order Issue DIQ Structure

FIG. 30 shows an enhanced DIQ 400 structure of the same floating-point ALU's DIQ 300 example in FIG. 27 to allow full out-of-order issue, although a DIQ for any functional unit could be used. Note, as shown in FIG. 30 (and FIG. 27), some of fields 485 (385 in FIG. 27) vary with the type of functional unit. The enhancements in the DIQ 400 comprise the additions of: (1)"issued" 486, "RS1_rdy" 491, and "RS2_rdy" 494 flags or fields in each entry (fields 487, 488, 489, 490, 492, and 493 are identical to fields 386, 387, 388, 389, 390, and 391, respectively, of FIG. 27), (2) comparators (not shown) to match a result tag with all operand tags (RS1 tag and RS2 tag) to update their ready flags (RS1 rdy and RS2_rdy), and (3) an issue logic 475 circuitry to determine which instruction entry should be issued next.

In the out-of-order issue DIQ 400 structure, an entry is still allocated from the tail side. Multiple tail pointers 460 are provided to handle multiple entry allocations in entry cells **495** per cycle. A newly allocated entry has its RS1_rdy 491 and RS2_rdy 493 fields initially set based on the operand value availability at decode time. These flags are updated during the writeback stage, by forwarding result tags 420 (NOT the result values as in the RS technique) to the appropriate functional units (for example, tags of floating-point results go only to selected functional units that use floating-point operands). These result tags 420 are compared to each entry's operand tags. A match will set the corresponding ready flag (RS1_rdy 491 or RS2_rdy 493) to TRUE.

Issue logic 475 selects an unissued instruction entry with both operands ready (RS1_rdy 491=RS2_rdy 493=TRUE). If there are multiple ready entries than the oldest one (the one closest to head pointer 480) is selected. The entry's issued flag is then set to TRUE. This entry will be popped from DIQ 400 when it reaches the bottom of DIQ 400. This simplification is necessary to reduce the complexity of issue logic 475 and allocate logic (not shown, but can be the same as allocate logic 345 in FIG. 29). If an issued entry is immediately freed then, a new entry can potentially be 50 allocated in the middle of DIQ 400's "queue". Consequently, the age or order of instructions inside DIQ 400 is no longer easily determined by issue logic 475. Multiple entry allocations from allocation ports 410 also become more complicated since they are not sequentially pushed at the tail side, but rather at random locations with possible contentions to resolve. Note that in contrast to DIQ 300 which issues all field values of fields 385 of a single cell **395** in-order during the issue operation to instruction issue register 516 (see FIG. 33), DIQ 400 issues all field 485 values during the issue operation (field values for fields 487, 488, 489, 490, 492, and 493 of a single cell 495) out-of-order to instruction issue register 516, except newly added field values 486, 491, and 494, which are discarded.

Combining MRB and DIQ

FIG. 31 shows superscalar processor 500 that utilizes DIQ 300 or 400 instruction shelving and MRB 100R or 100F

(both have identical structure to MRB 100, but MRB 100R buffers fixed-point register values, while MRB 100F buffers floating-point register values) result shelving techniques. Processor 500 has execution units 560, 570, 580, 590, 595, 596, 505, and 506 analogous to the execution units shown in FIG. 18. Superscalar processor 500 as shown in FIG. 31 clearly reveals a significant reduction in global buses and multiplexers compared to the processor of FIG. 18. By eliminating operand value copying and result value forwarding, shared-global operand buses and both sharedglobal result buses and wide-bus multiplexers are avoided, replaced by private-local (module-to-module) read buses and write buses 525, respectively. The only shared-global buses left are the required instruction dispatch buses 550 to deliver decoded instructions to every execution unit's DIO 300 or 400. In the case of out-of-order issue DIQs 400 a small number of global wires to carry result tags are added (not shown in FIG. 31).

In processor 500, separate register units 530 and 540 are provided for fixed-point and floating-point register results. 20 (It is also possible to combine both types of register results in a single register unit.) With two register units, the processor 500 core area is basically segmented by Fixed-Point Register Unit (FXRU)530 to hold general-purpose "R" registers, and the Floating-Point Register Unit (FPRU)540 25 to hold floating-point "F" registers. Special purpose registers for condition codes can use any of the "R" registers, following the "no-single copy of any resource" philosophy of the DEC ALPHA architecture. A single-copy of any resource can become a point of resource contention. The "R" and "F" registers (which may be contained in either register units 530 and 540 in 515R and 515F or in MRB 100R and MRB 100F) are also used to hold fixed- and floating-point exception conditions and status/control information, respectively.

Each of register units 530 and 540 contain register files 515R and 515F accompanied by MRBs 100R and 100F, respectively, to support register renaming, out-of-order execution, multi-level speculative execution, and precise interrupts. FIG. 32 shows the organization of a FXRU 530 40 (FPRU 540 is similar). RF 515R contains the in-order state, while MRB 100R contains the look-ahead state. An MRB 100R entry is retired to RF 515R only if it is safe. To read an operand, the register tag (reg_tag, a unique associative search key) is presented to MRB 100R and the register 45 number (reg_num) is presented to RF 515R, thus performing the read on both RF 515R and MRB 100R. If a match entry is found in MRB 100R (read_found=1), then the register content in RF 515R is considered old and ignored. RF 515R's read gives the correct register value. Note that finding a matched entry in MRB 100R does not guarantee that the register value has been found. The register value may still be computed in one of the execution units (560, 570, 580, 590, 595, 596, or 506) and has not been written to 55 MRB 100R. MRB 100R includes allocate ports 110R, write ports 120R, read ports 130R, and retire ports 140R similar to MRB 100.

Fixed-point arithmetic instructions are executed in the Fixed-Point Units (FXU 0 and FXU 1)570 and 580. 60 Floating-point arithmetic instructions are executed in the Floating-Point Arithmetic Logic Unit (FALU)506 and Floating-Point Multiply/Divide/Square-Root Unit (FMDS) 505. Note that FALU 506 also performs floating-point compare/set instructions and writes its condition code 65 directly to FXRU 530. Conditional/Immediate Move Unit (CIMU)596 performs register move instructions between

FXRU 530 and FPRU 540, as well as within FXRU 530/ FPRU 540 itself. CIMU 596 can also be dedicated to handle conditional move instructions as seen in the SPARC-64 (version 9) instruction set. Load & Store Units (LSU 0 and LSU 1)590 and 595 perform all load and store operations and include store queues (SQ)591 and 594, respectively, to queue the store instructions until they can be committed/ executed safely, with load bypassing and two simultaneous data accesses to D-cache 511 allowed. It is to be understood that one or more memories or other storage units may be employed instead of a cache for D-cache 511 for performing an equivalent function as D-cache 511. Branch instructions and PC address calculations are executed in the Instruction Address Unit (IAU)560. A BTB (branch target buffer), 15 which is a combination of a branch-target address cache (or other memory or storage unit) and branch history table, is provided in IAU 560 to help eliminate some branch delays and predict branch direction dynamically. During processor 500 implementation, it is best to physically layout circuit modules/blocks such that execution units 560, 570, 580, 590, 595, 596, 505, and 506 surround their corresponding register unit 530 or 540. Execution units that access both register units 530 and 540, such as LSUs 590 and 595 and CIMU 596, can be placed between the two. In this way, local bus 525 wiring is more direct and shorter.

DIQ 300 or 400 resides in each execution unit (560, 570, 580, 590, 595, 596, 505, and 506). FIG. 33 shows FALU 506 organization, which gives an example how DIQ 300 or 400 is connected. An instruction issue register can be added to 30 reduce the operation delays in the issue/execute pipeline. To allow immediate instruction execution following an instruction allocation to an empty DIQ 300 or 400, a bypass route from the instruction dispatch buses 550 to instruction issue register 516 is provided through multiplexer 517.

During normal operations, arithmetic and load/store instructions proceed through five processing steps/stages; fetch, decode, issue/execute, writeback, and retire. (Note that a stage does not necessarily represent a single hardware pipeline stage that can be performed in one clock cycle.) At the fetch stage, multiple instructions (N_d) are fetched through fetch buffer **521** simultaneously from I-cache arrays (or other memory or storage unit)510 (see FIG. 34). With instruction aligning done in I-cache arrays (or other memory or storage unit)510 (as in the IBM RS/6000 processor-see Grohoski), N_d instructions can be fetched each cycle, without wasted slots, as long as they reside in the same cache (or other memory or storage unit)510 line. At the decode stage, multiple fetched instructions are decoded by instruction decoder 524 simultaneously. Each valid instruction is However, if the register is not found in MRB 100R then the 50 assigned a tag by IIU (Instruction ID Unit)522, which is also used to tag the destination register of the instructions. An entry is allocated at the "R" (100R) or "F" (100F) MRB for each new register assignment. Register tags of an instruction's operands are acquired from the RTU (Register Tag Unit)523. Finally, at least one valid decoded instruction or all (or multiple) valid decoded instructions are dispatched to the appropriate execution unit's (560, 570, 580, 590, 595, 596, 505, or 506) DIQ 300 or 400. Decoded instructions are shelved by DIQs 300 or 400 to allow more time to resolve data dependencies. Each DIO 300 or 400 includes dependency check logic that automatically issues an instruction at the bottom of DIQ 300 or 400 as soon as its operands become available, independent of other DIQs 300 or 400. Unlike typical von Neumann processors, no centralized control unit is required to explicitly and rigidly sequence every instruction, deciding when it can execute. This is the essence of dynamic, dataflow scheduling. At burst

60

65

situations, all execution units 560, 570, 580, 590, 595, 596, 505, and 506 simultaneously issue an instruction, achieving maximum machine parallelism. Results are not directly written back to their register file 515R or 515F, but shelved first at MRB 100R or 100F. Retiring of an instruction's result from MRB 100R or 100F to register file 515R or 515F is done when safe, i.e., (a) there is no exception in the execution of the instruction and instructions preceding it, and (b) there is no prior conditional branch instruction that is outstanding or unresolved. This ensures correct execution of a program, giving the same results as if the program was run sequentially. Retiring of a store instruction, which involves a permanent write to D-cache (or memory, I/O device, or other storage unit)511, follows the same procedure. A summary of the flow of operations involved in each processing stage is depicted in a flowchart shown in FIGS. 35 and 36.

A unique retire process, using the branch_point and in_order_point, has been introduced. The branch_point (generated by IAU 560) is the inst_ID of the "oldest" 20 unresolved conditional branch (inst_ID of IAU 560's bottom DIQ 300 or 400 entry). Therefore, all instructions prior to the conditional branch (inst_ID<branch_point) are nonspeculative. The in_order_point (generated by IIU 522) is the inst_ID of the "oldest" instruction that has not completed or completed with an exception. Thus, if an instruc- ²⁵ tion has inst_ID<in_order_point, then its preceding instructions completed without an exception. Unlike Johnson's RB and the Metaflow THUNDER SPARC's central DCAF, this retire process using the branch point and in_order_point eliminates the need for allocating "dummy 30 entries" to the result shelf for branches, stores, or any other instructions that do not write to a register file. It also eases synchronization in multiple result shelves and store buffers.

Cost and Performance Analysis

In Sections 5 and 6, respectively, of Chamdani 35 dissertation, the subject of which is incorporated herein by reference, cost and performance analyses are addressed for MRB 100 and DIQ 300. The cost DIQ 400 is slightly higher in number of comparators and global wires, however performance is also increased slightly. For MRB 100, the cost 40 analysis suggests a major reduction in the MRB 100's usage of multiplexers and global wires, relative to the RB's usage. MRB 100's cost complexity/growth for multiplexers is $O(n^2)$ while the RB's is $O(n^3)$ where O is interpreted to mean the order-of-magnitude of what is contained between 45 parentheses and n is the number of individual units of the particular item being referred to. The complexity of existing RB techniques (Unlike the RB, the number of global wires used via the MRB 100 is unaffected by the data size N_{data}) is a proven fact. AMD found out in their superscalar 29K 50 processor (see Case), an "intellectual predecessor" to the K5, that a mere 10-entry RB plus its operand and result buses consumed a significant chip area, about the same as 4-Kbyte cache! And this is with the chip fabricated in an advanced 0.5-micron CMOS technology.

The performance analysis confirms that the MRB 100 technique does work, and retains all of RB's good performance characteristics. Based on the benchmarking done on a 4-way superscalar model, the performance was observed to saturate at a relatively low number of 12 to 16 MRB 100 entries. With dummy branch entry allocation in the RB for conditional branches, the RB's performance is expected to saturate at a higher number. Or, if the number of entries in the RB is set based on MRB 100's saturation number, then the RB's performance could be lower due to the additional decode stalls from RB full.

For DIQ 300 or 400, the cost analysis suggests an improvement in almost every hardware component, with

major reductions in the use of comparators, multiplexers, and global wires with respect to the reservation station technique. The expensive shared-global wires are mostly replaced by private-local wires **525** that are easier to route, have less propagation delay, and occupy much smaller silicon area. DIQ **300**'s or **400**'s number of global wires remain the same as the number of DIQ **300** or **400** entries and data size increase. DIQ **300**'s or **400**'s cost complexity/ growth for comparators is $O(nlog_2n)$ or $O(n^2)$ compared to $O(n^3)$ for RS. For multiplexers, DIQ **300** cost grows as $O(log_2n)$ while the RS cost grows as $O(n^2)$. Also, DIQ **300**'s or **400**'s number of comparators is not affected by increases in entry size or data size. This makes the DIQ **300** or **400** technique more adaptable to future 64-bit superscalar architectures.

On the performance side, the good characteristics of the RS technique in achieving maximum machine parallelism have been maintained in the DIQ 300 or 400 technique. The only sacrifice made in DIQ 300 technique is the use of in-order issue with an instruction window. This may penalize performance slightly on the cycle count, which can be easily recovered through faster and simpler circuit implementation. In the end, the actual speed or performance of the processor is faster due to reduced cycle time or more operations executed per cycle. (The out-of-order issue DIQ 400 technique is at par with the RS technique in terms of cycle-count performance, but higher in terms of overall performance if the improved clock frequency is factored in.) The performance analysis confirms that a good performance speedup, on the cycle count basis, is still achieved. Based on the benchmark set used, a speedup between $2.6 \times \text{to } 3.3 \times \text{was}$ realized in a 4-way superscalar model over its scalar counterpart. Moreover, the performance saturates at a relatively low number of 4 DIQ 300 or 400 entries. These results can be compared to 4-way superscalar processors which typically gain less than 2.0× over scalar designs on the SPE-Cint92 benchmarks (see L. Gwennap, "Architects Debate VLIW, Single Chip MP," Microprocessor Report, Vol. 8, No. 16, pp. 20-21, Dec. 5, 1994).

The following sections give a more detailed description of the IIU **522** and RTU **523** blocks in Fetch & Decode Unit (FDU)**520** shown in FIG. **31**.

Inst_ID Unit (IIU)

Inst_ID Unit (IIU)522 has several responsibilities; (a) assign inst_ID to every decoded instruction, (b) record the PC address of every decoded instruction, (c) give the PC address of an exception instruction, and (d) determine the oldest instruction that has not completed or completed with an exception (its inst_ID is called in_order_point). The main component of IIU 522 is an Instruction Array (IA), which stores the PC address and status of every active instruction currently in a processor (for example, processor 500). The IA is implemented as a circular FIFO buffer with head and tail pointers (head_IA and tail_IA). During normal instruction run, the next tail_IA is incremented based on the number of valid instructions in the decode group. However, when a branch misprediction is detected, the tail_IA must be adjusted. All IA 522 entries following the mispredicted branch entry are flushed. Thus, next tail_IA is set to the IA 522 address of the mispredicted branch entry, which is the branch_point less its color_bit. Note that the branch_point (generated by IAU 522) is the inst_ID of the current "oldest" unresolved conditional branch.

A tag or active instruction identifier, inst_ID, is composed of (color_bit, IA_addr), where IA_addr is the address/index of an entry in the IA (color_bit is defined below). It provides the "age/order" information during instruction-order comparisons. An instruction with a smaller inst_ID indicates that the instruction is "older", because it
15

30

50

55

60

appeared earlier in the decode sequence. NULL inst_IDS with IA addr=0, (0,0) and (1,0), are reserved to indicate non-existence or no-value. This is important to indicate no pending conditional branch instruction, branch_point=(0,0) or (1,0). It can also mean no register operand (e.g., replaced by immediate value, does not have second operand, default zero register value in R0, etc.) when register tag is (0,0) or (1,0).

The color_bit (as used in the Metaflow architecture-see Popescu, et al.) is a single binary bit that is appended to distinguish instruction order when the IA address starts to wrap around. It is toggled at each wrap around. For example, consider an inst_ID with 7-bit IA_addr (array size of 128). If an instruction has inst_ID of (0,127), its next decoded instruction has inst_ID of (1,1). When the two inst_IDs are compared, (0,127) < (1,1). In the next wrap around (now the color_bit is toggled from 1 to 0), (1,127)<(0,1). The rule for comparing two inst_IDs (inst_ID_A and inst_ID_B) is:

if (color_	_bit _A	≠	color_	_bit _B)	then

/* reverse compare result * compare(inst_IDA,inst_IDB)=NOT compare(IA_addrA,IA_addrB) else

* same color bit */

Thus the sense of comparison is reversed when the color bits are different.

Another responsibility of IIU 522 is to determine the in_order_point, the inst_ID of oldest instruction that has not completed or completed with an exception. The in_order_point will be used during instruction retirings. The in_order_point less its color_bit is actually the IA address of oldest IA entry with "complete" bit=0 (the one closest to head_IA). The ideal way of computing in_order_point is to check the "complete" bits of all IA entries between head_IA and tail_IA. However, this is not efficient for hardware implementation. Instead, the checking 40 can be limited to within a window of N_{iw} entries. This "in_order_point" window starts from the IA address of current in_order_point (called iop_IA_addr) to (iop_ IA_addr+N_{iw}) mod N_{ia}. This simplification will not alter program correctness. The only potential penalty is a postponement in instruction retiring, if the in_order_point 45 moves "up" (toward tail_IA) too slow. This should be infrequent if the size of "in_order_point" window (N_{iw}) is not too small.

Register Tag Unit

FIG. 37 shows the organization of Register Tag Unit (RTU) 523. RTU 523 maintains the most recent tag of every "R" and "F" register (which may be in 515R and 515F or in MRBs 100R and 100F). The most recent tag of a register is the inst_ID of producer instruction that updates the register last. To store tags of all "R" and "F" registers (may be in 515R and 515F or in MRB 100R and 100F), Register Tag File (RTF)600 is used. RTF 600 has the structure of a register file, except the register content is not a value, but rather a tag. To support speculative execution, Register Tag Reorder Buffer (RTRB) 100RT accompanies RTF 600. RTF 600 are similar to 515R or 515F except that it holds to register values, only register tags, and it was in_order_ point to clear the stale/old register tag. (Note that RTRB is not similar to MRB 100.) RTRB 100RT has RTU allocate ports 810, both RTRB 100RT and RTF 600 share read ports 65 830, and tags are updated from RTRB 100RT to RTF 600 through RTRB retire ports 840. Register tags of source

operands RS1(0. N_d -1) and RS2(0. N_d -1) are read from both RTF **600** and RTRB **100**RT (see left side of FIG. **37**). If a match is found in RTRB, then it contains the most updated register tag, and the tag in RTF is considered old. Tags in RTRB 100RT are updated to RTF 600 as quickly as possible, provided that all previous conditional branches have been resolved. Thus, RTRB 100RT will mostly contain 'speculative" register tags. These tags are flushed when the conditional branch is found to be mispredicted. To keep up with the instruction fetch rates up to N_d entries at the bottom of RTRB 100RT can be simultaneously retired to RTF 600. FIG. 38 shows the structure of RTF 600. RTF cells 695 which include DFFs 605 (or equivalents thereof) are shown in FIG. 38. Also shown in cell 695 is comparison unit 652 (may be an equivalent device to perform the same function as would be understood to those of ordinary skill in the art) for clearing the stale/old register tag. Note that Since there can be multiple updates to the same register location among N_d retirees, a priority selector is accommodated to make sure

only the last update is written. Also note that a difference between RTRB 100RT and RTF 600 with regard to speculative tags is that RTRB 100RT holds speculative register tags while RTF 600 holds nonspeculative tags of active instructions (not stale).

RTRB 100RT is implemented as a multi-ported, content-35 addressable, circular FIFO buffer with the register number (reg_num 730—see below) used as the associative search key during read operations. The circular FIFO control is implemented by a set of head and tail pointers (head_RTRB and tail_RTRB) which are not shown, but are analogous to head pointers 180 and tail pointers 160 of FIG. 21. The content of an RTRB entry 700 for RTRB 100RT is shown in FIG. 39 with fields n_empty 710, "last" 720, reg_num 730, and reg_tag 740. To avoid prioritized associative lookup, "last" bit 720 is provided for each RTRB entry 700. Thus, the most recent tag (740) of a register number (730) is read from the matched RTRB entry 700 with "last" bit 720=1. "Last" bit 720 is set during entry allocation if its register number 730 is the most recent one within the decode group. Each entry's "last" bit 720 is reset when a more recent entry allocation "overwrites" its register number 730.

The structure of RTRB 100RT is shown in FIG. 40. Each RTRB cell 895 consists of D flip-flops (DFFs) or RS flip-flops (or other storage devices)805 to hold RTRB entry 700, and logic circuits for allocate 845, read 825, and retire 835 operations. RTRB 100RT is a simplified reorder buffer with the following modifications; (1) buffer only register tags 740, no register values or program counters, (2) no write port, (3) multiple allocate and retire ports, and (4) allow read bypassing from allocate ports.

The fourth modification, read bypassing, is necessary because the most recent tag 740 of a register may still be in allocate ports 810, not yet written to RTRB 100RT. Consider the following example where $((x-y)^2+z)^2$ computation is about to take place. Assume all calculations are performed in single-precision floating point arithmetics; variables x, y, z were already loaded into register F1, F2, F3, respectively; and the result is written to register F4. Suppose the current decode group is as follows: $(N_d=4)$

 $compare(inst_ID_A, inst_ID_B) = compare(IA_addr_A, IA_addr_B).$

15

i	inst_ID(i)	opcode(i)	RD(i)	RS1(i)	RS1_tag(i)	RS2(i)	RS2_tag(i)	Operation
0	(0,126)	sub.s	F4	F1	n/a	F2	n/a	F4 <- F1 - F2
1	(0,127)	mul.s	F4	F4	(0,126)	F4	(0,126)	F4 <- F4 * F4
2	(1,0)	add.s	F4	F4	(0,127)	F3	n/a	F4 <- F4 + F3
3	(1,1)	mul.s	F4	F4	(1,0)	F4	(1,0)	F4 <- F4 * F4

Consider the register tag reading of the third instruction (i=2) in the decode group. When RS1(2) is presented to RTRB cells 895, it is also checked for possible bypassing at the RS1 Bypass Logic unit or block 875 (see FIG. 40). Because allocations and reads occur in the same cycle, RTRB 100RT does not have the current decode group's register tags available for reading. The register tag read must be satisfied directly from one of allocate ports 810. In

where i $\in [1, N_d - 1]$. The increment from the base tail pointer is determined based on the number of allocations in the previous ports. Note that the next cycle tail_RTRB is simply $(tail_RTRB(N_d-1)+RTRB_alloc_en(N_d-1)) \mod N_{rtrb}$. The Tail Pointer Logic 850 also activates the RTRB_full flag if an allocation is about to overwrite the head entry. Thus,

$$RTRB_full = \begin{cases} 1, \text{ if } (head_RTRB \neq \text{tail}_{RB}) \cap \left(\bigcup_{k=1}^{N_d-1} (head_RTRB = ?tail_RTRB(k))) \\ 0, \text{ otherwise} \end{cases}$$
(11)

25

general, bypassing at RS1 read ports 830 are enabled with the following condition:

$$RSI_bypass_en(i) = \begin{cases} 1, & \text{if } \bigcup_{k=0}^{i-1} (RSI(i) = ? RD(k)) \\ 0, & \text{otherwise} \end{cases},$$
(9)

where i ϵ [1,N_d-1]. In this case, RS1_bypass_en(2)=1 because RS1(2)=RD(1) and/or RS1(2)=RD(0). Despite two match ups, the RS1_bypass(2) must come from the second allocate port (inst_ID (1)) of allocate ports 810 with inst_ ID=(0,127) because it updates F4 last. With RS1_bypass en(2)=1, the RS1_tag(2) is satisfied from RS1_bypass(2) 40 and not from RTRB 100RT. Read bypassing at RS2 ports (of read ports 830) follows the same analogy using RS2 Bypass Logic unit or block 865. obviously, RS1 bypass_en(0) and RS2_bypass_en(0) are always 0 because the register tags of the first instruction's source operands would have been updated to RTU 523 in the previous cycle(s).

The entry allocation at the tail (determined by tail pointer logic 850) of RTRB 100RT involves setting n_empty 710= 1, last 720=1 (if the register assignment is the last update— 50 controlled by last update logic 855-among instructions in the decode group), reg_num 730=RD, and reg_tag 740= inst_ID. There can be up to N_d RTRB 100RT entries allocated in a cycle. An allocation of an RTRB 100RT entry is enabled (RTRB_alloc_en(i)=1) if the fetched instruction ⁵⁵ at position i is not following a branch-taken instruction and writes to an "R" or "F" register (R_alloc_en(i) or F_alloc_en(i) is set), which may be in 515R and 515F, or in MRB 100R and MRB 100F. The Tail Pointer Logic unit or block 850 determines the correct tail pointer position for 60 every allocation port 810:

$$tail_RTRB(i) = \left(tail_RTRB + \sum_{k=0}^{i-1} RTRB_alloc_en(k)\right) \mod N_{rtrb},$$
(10)

35

If RTRB_full=1, the fetch and decode stage are stalled until there is room in RTRB 100RT.

The Last Update Logic block 855 determines whether a register assignment is the last update within the decode 30 group, so that the "last" bit can be set properly during allocation. For a register assignment to be the last update, there must be no subsequent assignment to the same register within the decode group. Thus, in general,

$$last_update(i) = \begin{cases} N_d^{-1} & (12) \\ 1, \text{ if } \bigcup_{k=i+1}^{N_d^{-1}} (RD(i) = ?RD(k)) \\ 0, \text{ otherwise} \end{cases},$$

where i $\in [0, N_d - 2]$. Obviously, last_update(N_d - 1) is always 1 because the last instruction (of the decode group) updates the last.

Finally, register tags 740 in the bottom RTRB cells/entries 895 are retired to RTF 600 if all its previous conditional branch instructions have been resolved. No "speculative" register tag may be updated to RTF 600. Thus, a retiree at retire_port(i) is green-lighted (valid_ret(i)=1) if its reg_ tag(i)<branch_point. The branch_point (generated by IAU 560) is the inst_ID of the current "oldest" unresolved conditional branch. Since all N_d entries at the bottom (cells/entries 895) are simultaneously evaluated to retire, head_RB(i)=(head_RB+i) mod N_{rtrb} is generated for $i=0 \dots N_d-1$. Each valid retirement moves the head pointer "up" by one position as determined by head pointer logic unit or block 870. So, the next cycle head pointer is,

(next)
$$head_RTRB = \left(head_RTRB + \sum_{i=0}^{N_d-1} valid_ret(i)\right) \mod N_{rirb}.$$
 (13)

Retiring is done as quickly as possible by providing N_{d} retire ports 840. Only "non-speculative" entries (reg. tag
branch_point) can be retired from RTRB 100RT to RTF 600. When a branch misprediction is detected, any "nonspeculative" entries left in RTRB 100RT are retired immediately, if any usually within a cycle. Then, RTRB 100RT is simply reset to repair from any misprediction side effects.

10

50

60

65

67

Another condition that causes flushing is when an exception error is detected. In this case, both RTRB **100**RT and RTF **600** are flushed. No register tags are needed when resuming from an exception, because all register values in FXRU **530**'s MRB **100**R and FPRU **540**'s MRB **10**OF would have been retired to RFs **515**R and **515**F, respectively, before going to the exception handler. To read register operands from RFs **515**R and **515**F, register number is used instead of register tag.

Finally, stale register tags in RTF **600** must also be cleared/reset to 0. The condition to reset is reg_tag <in_ order_point, which means that the reg_tag has been deal-located from IIU **522**.

Numerous and additional modifications and variations of the present invention are possible in light of the above teachings. It is therefore to be understood that within the scope of the appended claims, the invention may be practiced otherwise than as specifically claimed.

What is claimed is:

1. A distributed instruction queue comprising:

- at least two entry cells, each of said at least two entry cells being capable of holding a decoded instruction corresponding to at least one source operand, said at least one source operand for each decoded instruction being ²⁵ stored in storage means external to said distributed instruction queue, said at least one source operand not being copied into said distributed instruction queue;
- an allocate port connected to said at least two entry cells for allocation of decoded instructions to said entry cells³⁰ without reading the at least one source operand corresponding to each of the decoded instructions; and
- an issue port connected to said at least two entry cells for issuance of decoded instructions under logic control in any order from said entry cells to a functional unit external to said distributed instruction queue after such time as all of the source operands corresponding to the decoded instruction to be issued are available in the external storage means, the source operands being read from the external storage means at such time as the decoded instruction is issued, and wherein no result value is forwarded from the functional unit to said distributed instruction queue.

2. The distributed instruction queue as claimed in claim 1, further comprising a tail pointer logic unit to adjust a tail pointer to point to the one of said at least two entry cells into which the next decoded instruction should be allocated from said allocate port.

3. The distributed instruction queue as claimed in claim **1**, further comprising:

- means for determining the one of said at least two entry cells in which the decoded instruction that was earliest 55 allocated is held; and
- a head pointer logic unit to adjust a head pointer to point to the one of said at least two entry cells in which the decoded instruction that was earliest allocated is held.

4. The distributed instruction queue as claimed in claim 1, further comprising an issue pointer logic unit to adjust an issue pointer to point to the one of said at least two entry cells from which the next decoded instruction is to be issued through said issue port.

5. The distributed instruction queue as claimed in claim **1**, wherein no operand values are copied into said distributed

instruction queue and no result values are received by said distributed instruction queue.

6. The distributed instruction queue as claimed in claim 1, wherein said distributed instruction queue is operated independently of any other distributed instruction queue.

7. The distributed instruction queue as claimed in claim 1, wherein said decoded instructions are issued in-order.

8. The distributed instruction queue as claimed in claim **1**, wherein said decoded instructions are issued out-of-order.

9. The distributed instruction queue as claimed in claim **1**, wherein said decoded instructions are issued in some form of limited out-of-order issue.

10. The distributed instruction queue as claimed in claim
1, wherein the distributed instruction queue issues the decoded instructions to a functional unit for execution, the functional unit being external to the distributed instruction queue.

11. The distributed instruction queue as claimed in claim
 20 1, wherein the distributed instruction queue shelves said decoded instructions.

12. The distributed instruction queue as claimed in claim 1, wherein said distributed instruction queue is a circular array.

13. The distributed instruction queue as claimed in claim 1, further comprising additional allocate ports connected to said at least two entry cells for allocation of decoded instructions to said at least two entry cells.

14. A distributed instruction queue comprising:

- means for receiving at least two decoded instructions, each of the at least two decoded instructions corresponding to at least one source operand, said at least one source operand for each decoded instruction being stored in storage means external to said distributed instruction queue, said at least one source operand not being copied into said distributed instruction queue;
- means for determining when all of the source operands for each of the decoded instructions are present in the external storage means without reading the source operands; and
- means for issuing each of the decoded instructions in any order to a functional unit external to said distributed instruction queue after determination by said determining means that all of the operands for the decoded instruction to be issued are available in the external storage means, the source operands being read from the external storage means at such time as the decoded instruction is issued, and wherein no result value is forwarded from the functional unit to said distributed instruction queue.

15. A method for issuing instructions in a superscalar processor, comprising the steps of:

- allocating at least two decoded instructions, each of the at least two decoded instructions corresponding to at least one source operand;
- determining when all of the source operands for each of the decoded instructions are available without reading the source operands and without copying the source operands;
- issuing each of the decoded instructions in any order after all of the source operands for the decoded instruction to be issued are available; and

10

reading the source operands for the decoded instruction issued in said issuing step at such time as the decoded instruction is issued.

16. The method of claim 15 wherein said issuing step further comprises issuing each of said decoded instructions in-order.

17. The method of claim 15 wherein said issuing step further comprises issuing each of said decoded instructions out-of-order.

18. The method of claim 15 wherein said issuing step further comprises issuing each of said decoded instructions in some form of limited out-of-order issue.

19. The method of claim **15** wherein said allocating step further comprises allocating the decoded instructions with-

out allocating the source operands corresponding to the decoded instructions.

20. The method of claim **15** wherein said issuing step further comprises issuing the decoded instructions without issuing the source operands corresponding to the decoded instructions.

21. The method of claim **15** wherein said issuing step further comprises issuing each of said decoded instructions to a functional unit.

22. The method of claim 15, further comprising a step of shelving said decoded instructions, said shelving step being performed after said allocating step and before said determining step.

* * * * *

UNITED STATES PATENT AND TRADEMARK OFFICE CERTIFICATE OF CORRECTION

Page 1 of 3

PATENT NO. : 6,112,019 DATED : AUGUST 29, 2000 INVENTOR(S) : Chamdani, et al.

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

Column 2, line 22, change "demands" to --demand--.

Column 4, line 8, delete "[?]".

Column 4, line 36, change "factors" to --factor--.

Column 4, insert "Summary of the Invention" between lines 52 and 53.

Column 6, line 10, change "instruction" to --instructions--.

Column 7, lines 22-23, change "typically s supports" to --typically supports--.

Column 8, line 8, change "1970, M. Johnson" to --1970; M. Johnson--.

Column 17, line 49, change "technologies" to -- Technologies--.

Column 20, line 38, change "Tomasulo" to -- Tomasulo's--.

Column 20, line 39, change "(1-to-1" to --(1-to-1)--.

Column 23, align underscore of heading "Comparisons of Instruction Shelving Techniques".

UNITED STATES PATENT AND TRADEMARK OFFICE CERTIFICATE OF CORRECTION

Page 2 of 3

PATENT NO. : 6,112,019 DATED : August 29, 2000 INVENTOR(S) : Chamdani, et al.

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

Column 28, lines 47-48, change "Metaf [line 47] low [line 48]" to --Metaflow--. Column 28, line 49, change "Metaf low" to --Metaflow--. Column 29, line 8, change "lines. of" to --lines. Of --. Column 30, line 42, change "Sum" to --Sun--. Column 31, line 10, change "ans" to --and--. Column 32, line 16, change "respect b integer" to --respect to integer--. Column 34, line 50, insert "," after --Unix--. Column 34, line 53, change "PLACODE" to -- PALcode --. Column 36, line 3, change "PA -7100" to --PA-7100--. Column 37, line 7, change "access" to --accessed--. Column 38, line 50, change "performs" to --perform--. Column 45, line 62, delete "to". Column 48, line 19, delete the first appearance of "to". Column 48, line 42, change "factors" to --factor--. Column 49, line 32, change "condition" to --conditions--. Column 50, line 32, add "." after --lookup--. Column 53, line 11, change "Retire" to --retire--. Column 54, line 54, change "logic 255 for d determining" to --logic 255 for determining--.

UNITED STATES PATENT AND TRADEMARK OFFICE CERTIFICATE OF CORRECTION

PATENT NO. : 6,112,019 DATED : August 29, 2000 INVENTOR(S) : Chamdani, et al. Page 3 of 3

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

Column 57, line 45, insert "in" between --used-- and --MRB--. Column 58, line 42, change "than" to --then--. Column 60, line 57, change "unit's" to --units--. Column 61, line 38, insert "of" between --cost-- and --DIQ 400--. Column 63, line 61, change "600 are" to --600 is--. Column 64, line 16, change "Since" to --since--. Column 67, line 5, change "MRB 100F" to --MRB 100F--.

Signed and Sealed this

Twenty-ninth Day of May, 2001

Hickolas P. Solai

NICHOLAS P. GODICI Acting Director of the United States Patent and Trademark Office

Attesting Officer

Attest:

ficer Acting