

**Computer Architecture:
Fundamentals, Tradeoffs, Challenges**

**Chapter 4: Because of Pipelines
(out-of-order execution, branch prediction)**

Yale Patt

The University of Texas at Austin


Austin, Texas

Fall, 2022

Outline

- **Out-of-order Execution**
 - The Tomasulo Algorithm (the IBM 360/91)
 - Pipelining with a hiccup
 - An example
 - What was wrong with it (did not have precise exceptions)
 - How do we fix it (the Reorder Buffer)
 - Pipelining with a second hiccup

• Branch Prediction

- Other ways to deal with conditional branches
 - Predication (the THUMB IT instruction)
 - Multiple pipelines
 - Delayed branch
 - Combine branches → 
- Compile time branch prediction (BTFN, hint bits)
- Runtime (LT, 2bitCtr, Two-level, Hybrid, Perceptron, TAGE)
- The HEP (Burton Smith, Donelcor (~1978))


The Scoreboard Has A Problem

Many instructions are unnecessarily held up.

Traffic Analogy: Cars that are blocked

How ~~to~~ deal with this blocking?

Answer: Get them out of the way

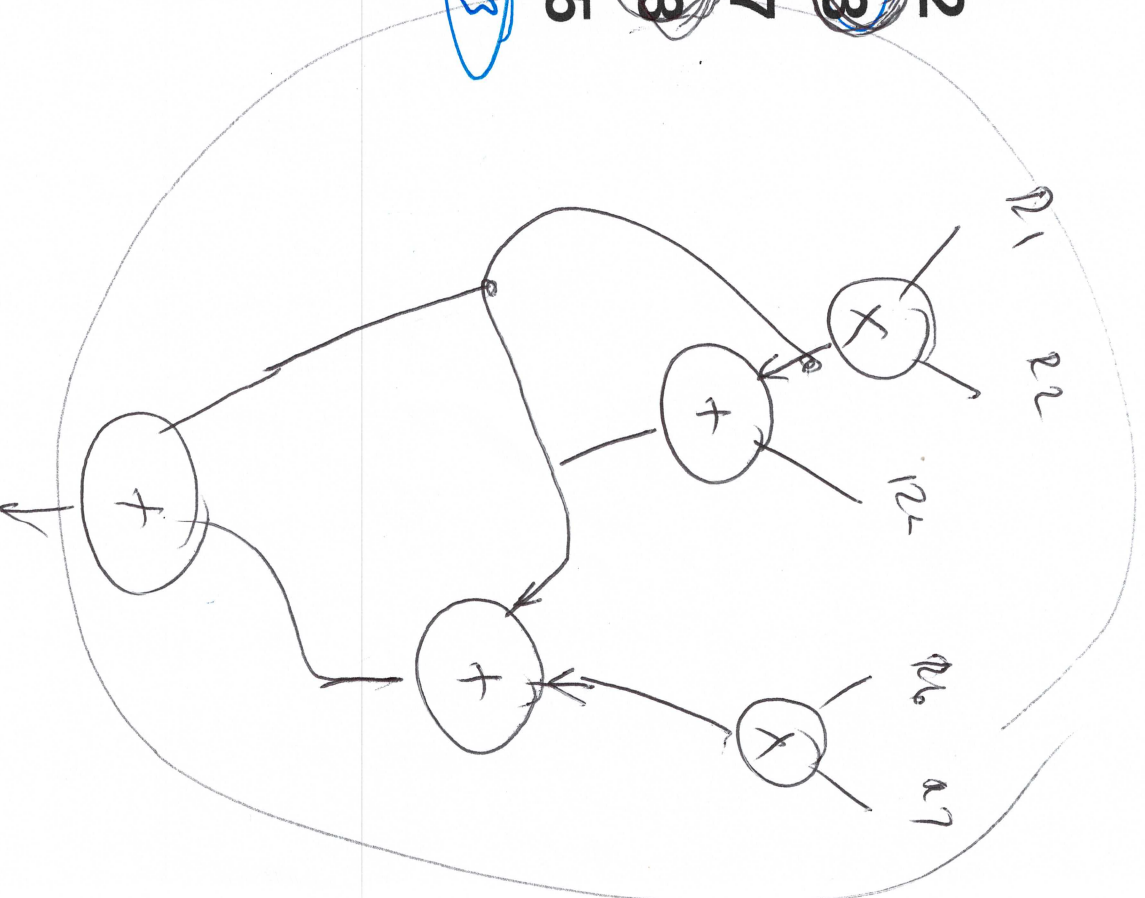


~~And a scoreboard buffer~~

An example

MUL R3, R1, R2
ADD R4, R2, R3
MUL R2, R6, R7
ADD R5, R2, R3
ADD R3, R3, R5

mul x, y r3



Out-of-Order execution (example)

- In-order

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	
MUL R3, R1, R2	F	D	✓	✓	✓	✓	(R3)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
ADD R4, R2, R3		F	D					✓	✓	(R4)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
MUL R2, R6, R7			F																					
ADD R5, R5, R3																								
ADD R3, R3, R5																								

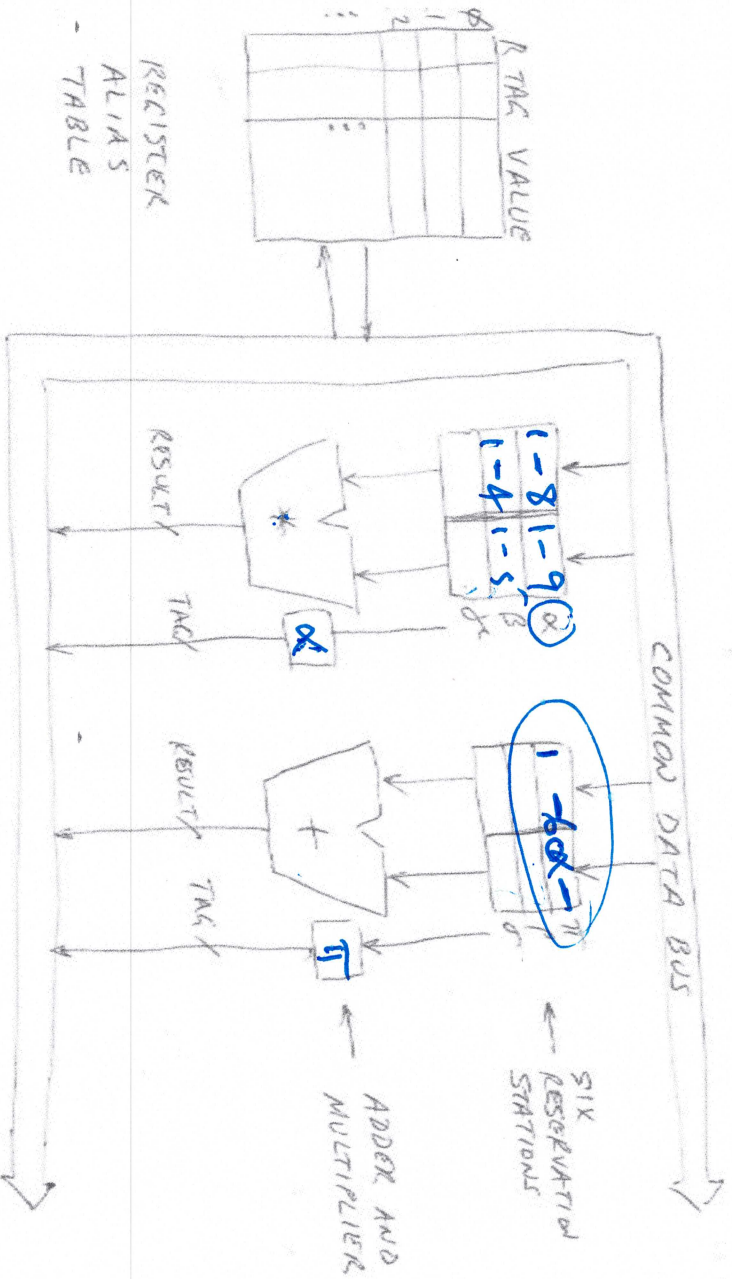

- Out-of-Order

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	
MUL R3, R1, R2	F	D	✓	✓	✓	✓	(R3)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
ADD R4, R2, R3		F	D																					
MUL R2, R4, R7			F																					
ADD R5, R2, R3																								
ADD R3, R3, R5																								





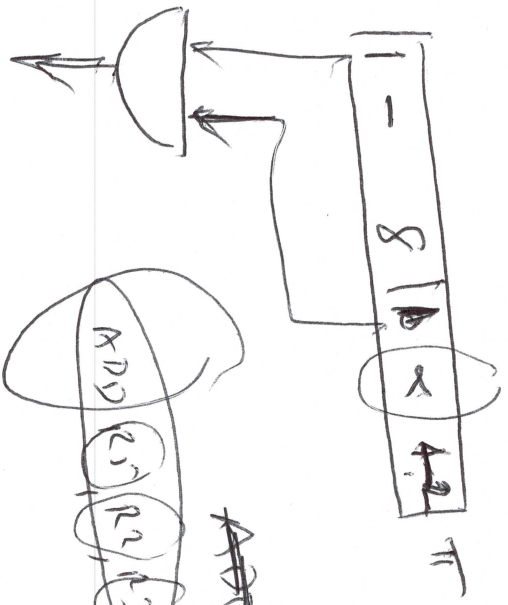

ADD R2, (R3), R7
 ADD R3, (R3), R5

ALIAS

Tomasulo Algorithm Microarchitecture




R ₀	1	1	6
R ₁	1	1	7
R ₂	1	X	8
R ₃	1		9
R ₄	1		10
R ₅	X		(10) 11
R ₆	1		12
R ₇	1		13



Out-of-order commit violates precise exceptions

- **Precise exceptions**
 - Identity of the faulting instruction is known
 - State recovered to just before faulting instruction
- **A simple example:**

DIV R1,R2,R3
ADD R2,R3,R4



ADD takes one cycle and stores result in R2


DIV eventually takes an exception.

Not possible to recover R2 prior to DIV execution

How To Fix The Problem (REORDER BUFFER)

- ADD A STRUCTURE (REORDER BUFFER)
- FETCH AN INSTRUCTION
- ALLOCATES AN ENTRY

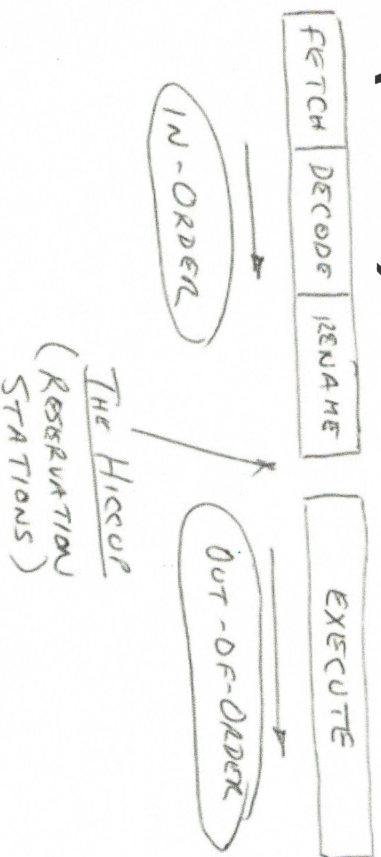
CODE
I₁: DIV R1, R2, R3
I₂: ADD R2, R3, R4
⋮



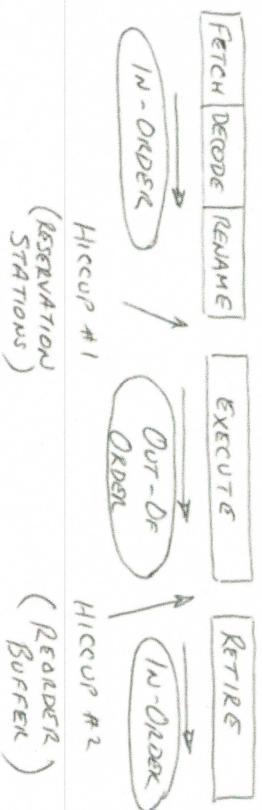
- FOR EACH INSTRUCTION, IS TEMPORARILY STORED
- ① RESULT OF EXECUTION (NOT IN THE REGISTER IN THIS REORDER BUFFER) RESULTS ARE STORED IN THE REGISTER.
 - ② ~~THE~~ ~~REGISTER~~ ALL OTHER REGISTER RESULTS ARE STORED IN THE REGISTER.

Pipelining with "hiccups"

- Tomasulo (~1965)



- HPS (~1985)

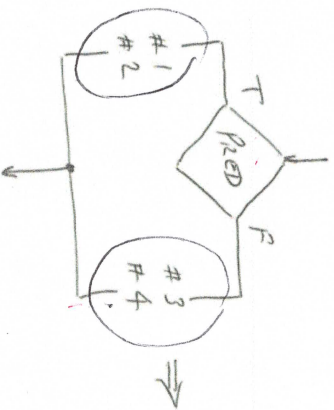


Branch Prediction

- **Ways to deal with Conditional Branches**
 - **Compile time**
 - **BTFN**
 - **Hint bits**
 - **Predication (the THUMB IT instruction)**
 - **Delayed branch**
 - **Combine branches**
 - **Run time Branch Prediction**
 - **LT**
 - **2bit Counter**
 - **Two-level**
 - **Hybrid**
 - **Perceptron**
 - **TAGE**
- **The HEP (Burton Smith, Donelcor (~1978))**

Predicated Execution (THUMB ISA)

- **Example 1**
(4 instructions)




↓
#1 (T)
#2 (T)
#3 (F)
#4 (F)

← 16 →

IT	P	16-bit
	#1	
	#2	
	#3	
	#4	

- **Example 2**
(3 instructions)




↓
#1 (T)
#2 (T)
#3 (F)

← 16 →



IT	P	16-bit
	#1	
	#2	
	#3	

Compiler Time



Run-Time

- LT
- 2-bit Counter




ISCA
1881
JIM SMITH

1 1 1 1 1 1 1 1 1 1



$$\frac{n-1}{n-2}$$




Dynamic branch prediction


- **Last time taken**
(Store a bit in the I Cache Tag Store)



- **Saturating 2 bit counter**
(Use high bit to predict)




- **Two-level Predictor**




- **Gshare Modification**

SCOTT McFARLAND



Two-Level Adaptive Branch Prediction The Basic Idea

1	0	0	1
0	0	0	1
1	0	0	1
1	0	0	1



TSB - Yo YEH

MICRO
(1951)



1/2V Partition Pro
TAGS - ANDRE SERENC
BRANCHLIST - SAVASBI ZAKHAROV

1101101101

1110101011110111

n

BRANCH HISTORY TABLE



PATTERNS
HISTORY TABLES

1111


97%

SCOTT McFALLINE
1993

KIMMING SO

IBM

ASPROS
1992



→

←

QSIFABLES

The Heterogeneous Element Processor (HEP)

- Burton Smith (at Donelcor ~1975)
- Introduction of Multithreading
 - Burton always credited CDC6600; I credit Burton!
- The mechanism (modified for explanation)
 - Four stages, HEP had 8
 - Four threads in progress, HEP had 8
 - Not shown: the hopper (20 threads waiting to participate)
 - *If a thread's instruction could not finish in time for next instruction*

cycle1 cycle2 cycle3 cycle4 cycle5 cycle6 cycle7 cycle8 cycle9

fetch1	decod	execu	store	fetch2	decod	execu	store	fetch3
fetch1	decod	execu	store	fetch2	decod	execu	store	
fetch1	decod	execu	store	fetch2	decod	execu		
fetch1	decod	execu	store	fetch2	decod			

Xie xie!