

# Out of Order Execution

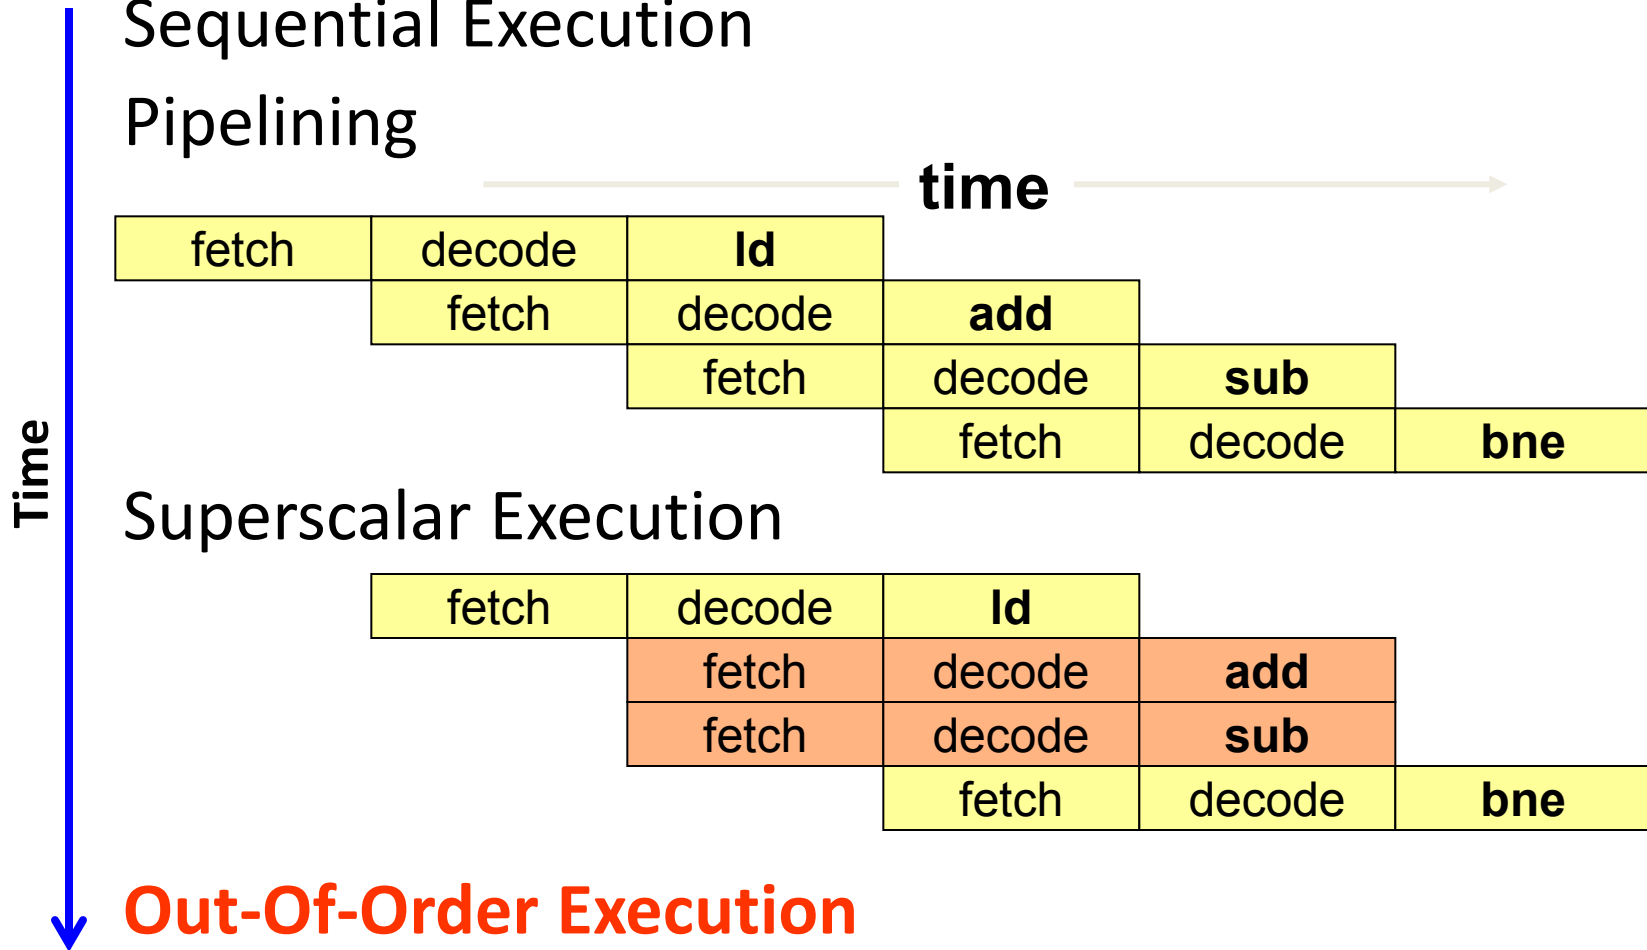
Pooja Saraff

Manoj Mardithaya

# How the notion was conceived


Sequential Execution

Pipelining



# What *is* OOO?

- OOO execution is a type of processing where the instructions can begin execution as soon as operands are ready
- Instructions are issued in order however execution proceeds out of order
- **Evolution**



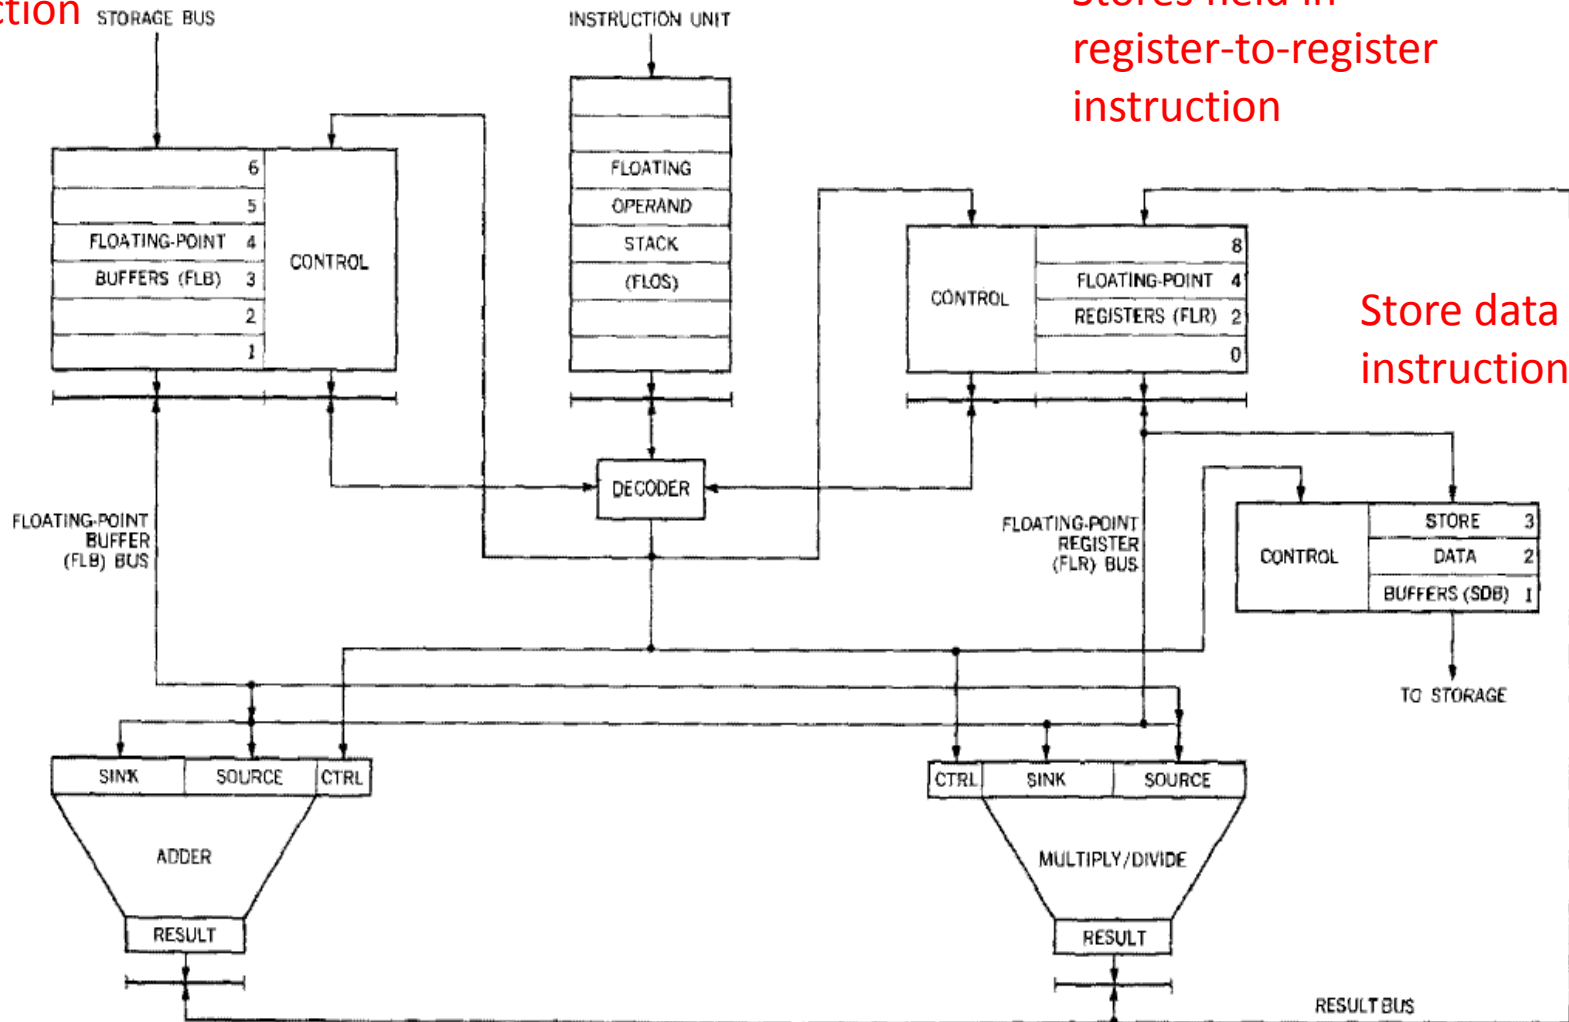
1964	CDC 6600
1966	IBM 360/91 <b>Tomasulo's algorithm</b>
1993	IBM/Motorola PowerPC 601
1995	Fujitsu/HAL SPARC64, Intel Pentium Pro
1996	MIPS R10000, AMD K5
1998	DEC Alpha 21264
2011	Sandy Bridge

# Architecture without Common Data Bus

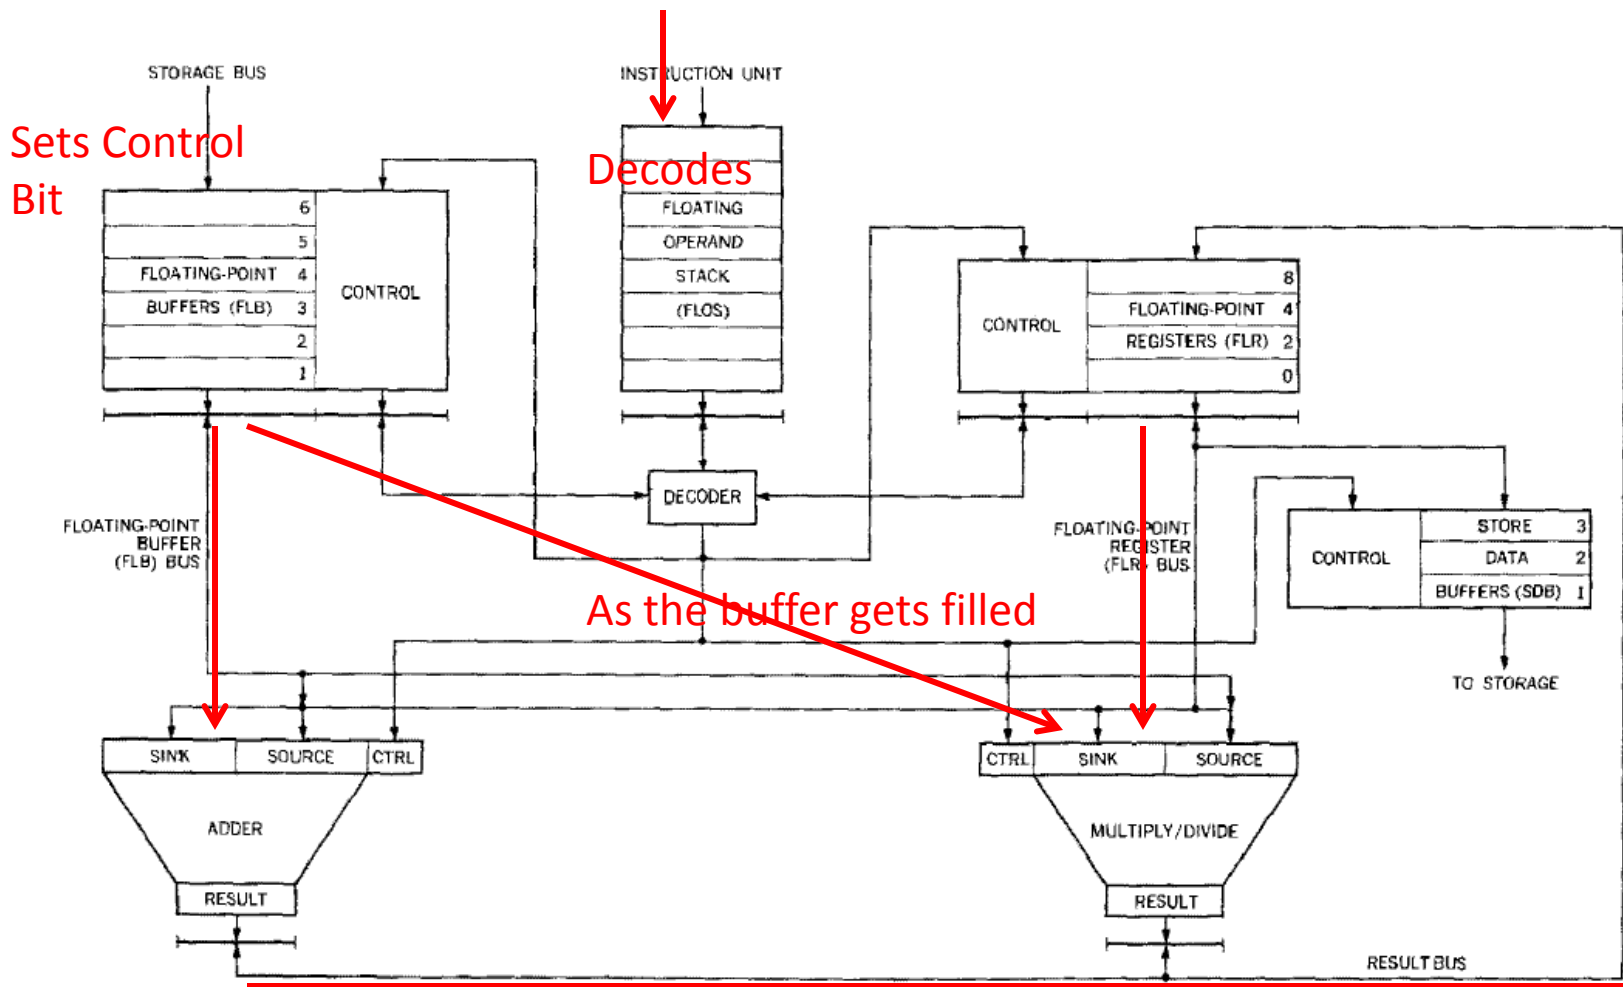
Storage-to-register instruction

Stores field in register-to-register instruction

Store data instruction



# Operation



Results

# Results

- It doesn't take care of data dependency
- Thus busy bit added – however FLOS hold-up because of busy sink register
- Solution to it – **Reservation Station**  
(control,sink,source)
- Execution now depends on appropriate reservation station

# 3 Types of Data Dependencies

- RAW (Read After Write)

$$\begin{aligned} R2 &\leftarrow R1 + R3 \\ R4 &\leftarrow R2 + R3 \end{aligned}$$

- WAR (Write After Read)

$$R4 \leftarrow R1 + R3$$

- WAW (Write After Write)

$$\begin{aligned} R3 &\leftarrow R1 + R2 \\ R2 &\leftarrow R4 + R7 \\ R2 &\leftarrow R1 + R2 \end{aligned}$$

Register Renaming uses in-order decoding to properly identify dependences.

# Register Renaming

<b>A: DIVF</b>	<b>F3,</b>	<b>F1,</b>	<b>F0</b>	<b>r1, -, -</b>
<b>B: SUBF</b>	<b>F2,</b>	<b>F1,</b>	<b>F0</b>	<b>r2, -, -</b>
<b>C: MULF</b>	<b>F0,</b>	<b>F2,</b>	<b>F4</b>	<b>r3, r2, -</b>
<b>D: SUBF</b>	<b>F6,</b>	<b>F2,</b>	<b>F3</b>	<b>r4, r2, r1</b>
<b>E: ADDF</b>	<b>F2,</b>	<b>F5,</b>	<b>F4</b>	<b>r5, -, -</b>
<b>F: ADDF</b>	<b>F0,</b>	<b>F0,</b>	<b>F2</b>	<b>r6, r3, r5</b>

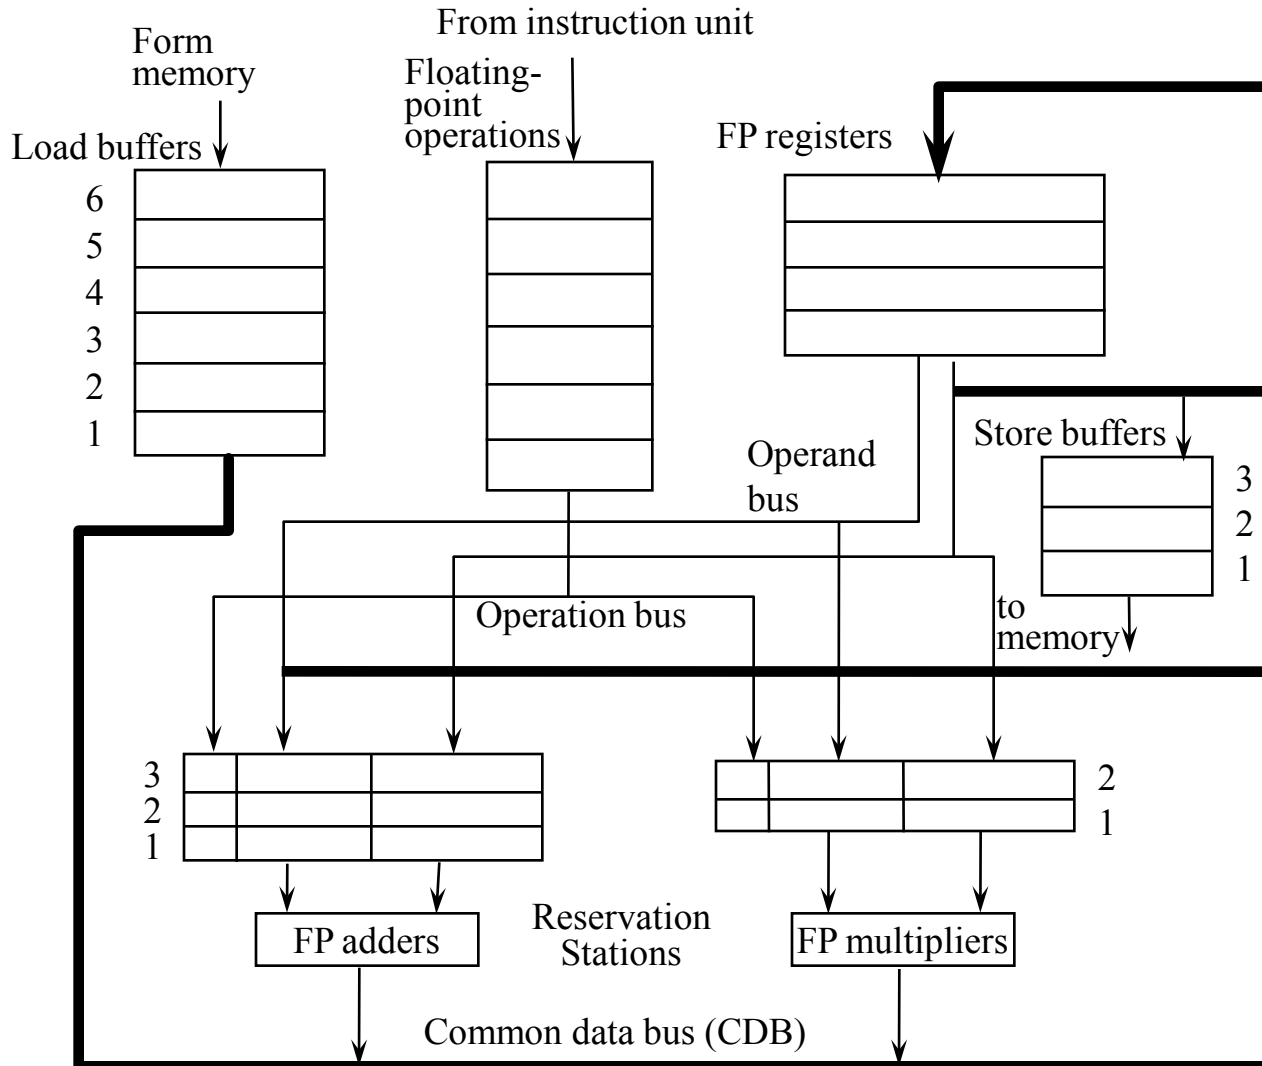
Register Rename Table

	<i>F0</i>	<i>F1</i>	<i>F2</i>	<i>F3</i>	<i>F5</i>	<i>F6</i>	<i>F7</i>	...	<i>F30</i>
<b>A</b>				R1					
<b>B</b>			R2	R1					
<b>C</b>	R3		R2	R1					
<b>D</b>	R3		R2	R1		R4			
<b>E</b>	R3		R5	R1		R4			
<b>F</b>	R6		R5	R1		R4			

Need more physical registers than architectural  
Ignores control flow for the time being.



# Architecture

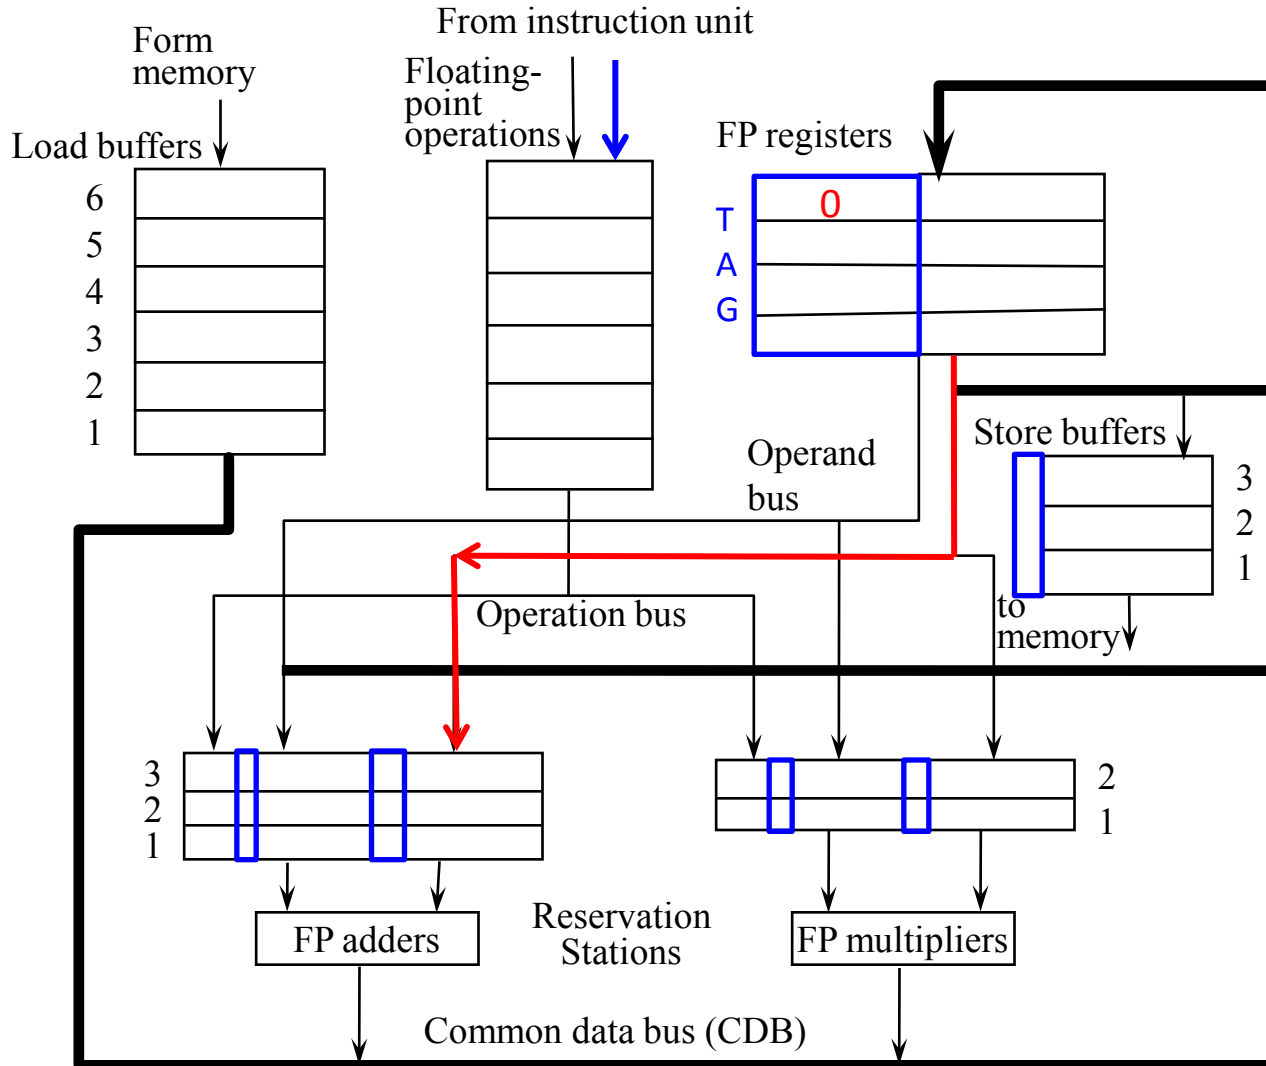


- 3 Adders
- 2 Multipliers
- Load buffers (6)
- Store buffers (3)
- FP Queue
- FP registers
- CDB: Common Data Bus

# Tomasulo's Algorithm Steps

- Issue
  - Issue if empty reservation station is found, fetch operands if they are in registers, otherwise assign a tag
  - If no empty reservation is found, stall and wait for one to get free
  - Renaming is performed here and WAW and WAR are resolved
- Execute
  - If operands are not ready, monitor the CDB for them
  - RAWs are resolved
  - When they are ready, execute the op in the FU
- Write Back
  - Send the results to CDB and update registers and the Store buffers
  - Store Buffers will write to memory during this step

# AD F0, FLB1

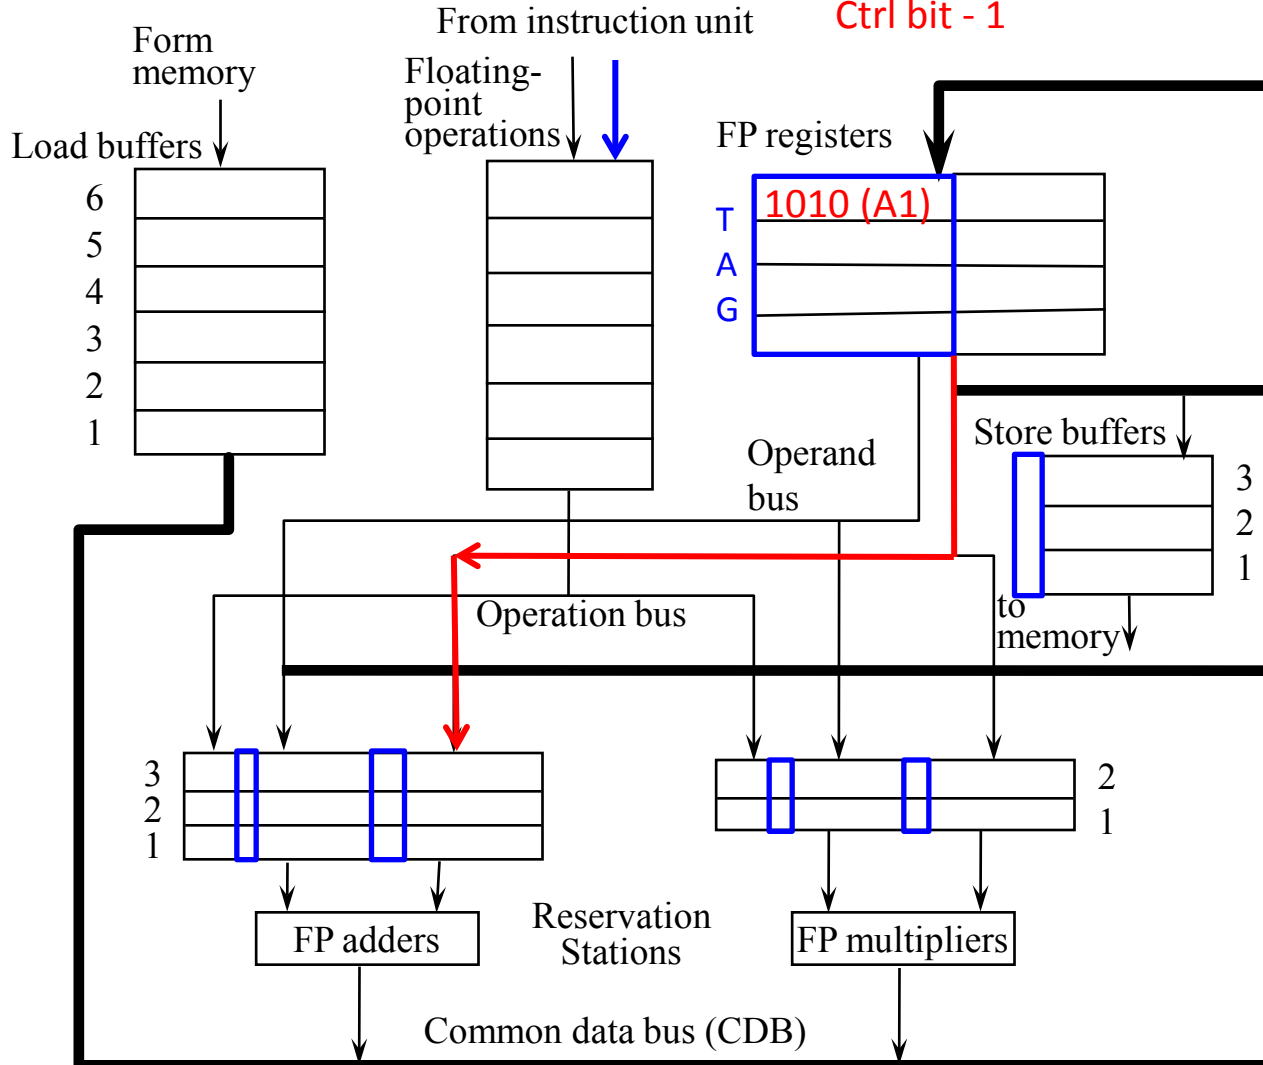


AD F0, FLB1

AD F0,.....

FLOS issues instruction and update Tag

Busy bit - 1  
Ctrl bit - 1



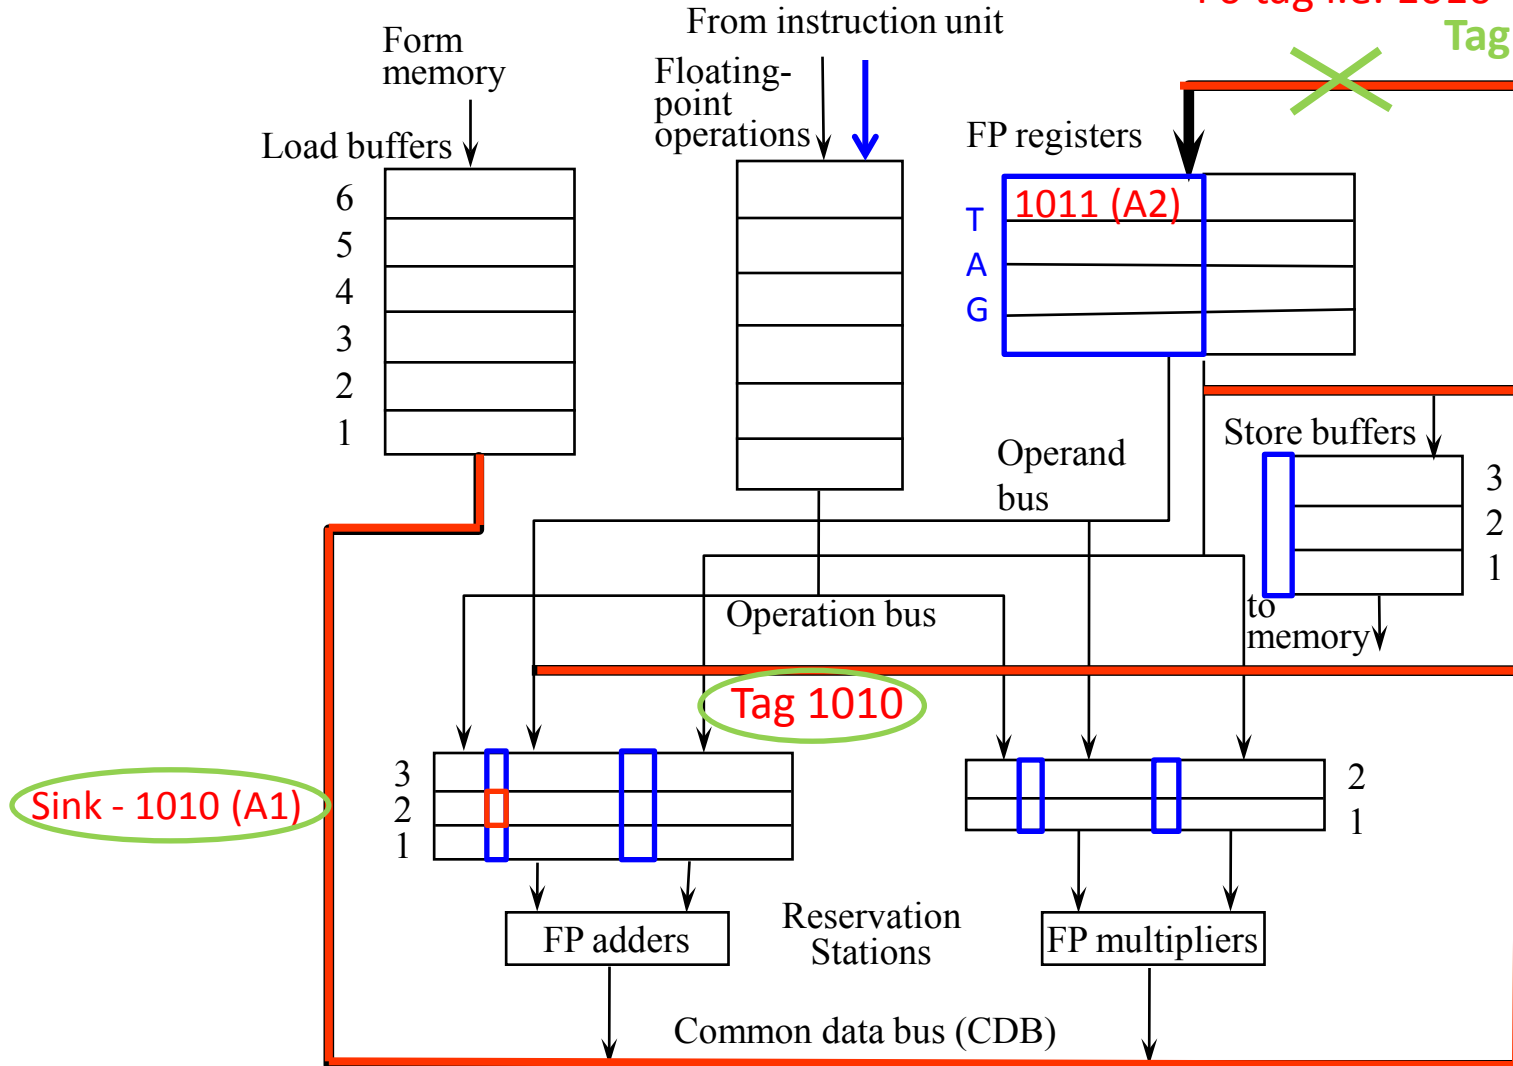
AD F0, FLB1 - execution

AD F0,.....

However reservation tag for A1 has initial

F0 tag i.e. 1010

Tag miss match



# Drawbacks

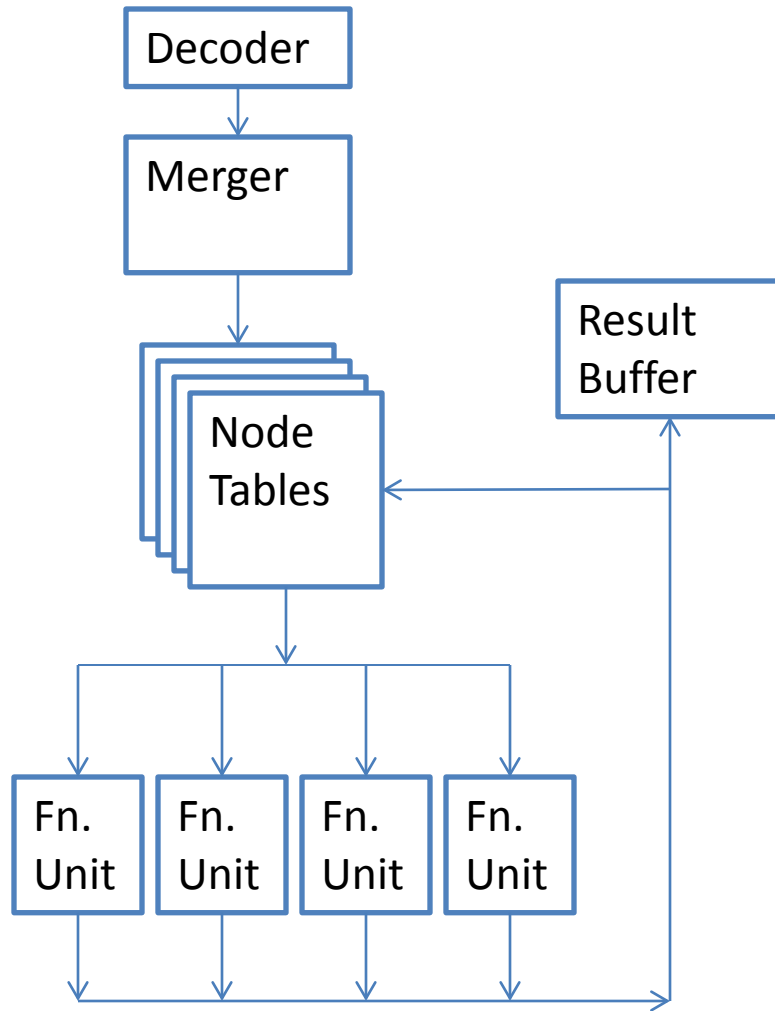
- CDB is a bottleneck
  - Limits the execution time of any instruction to 2 cycles, minimum
- Complex implementation

# HPS: a new microarchitecture

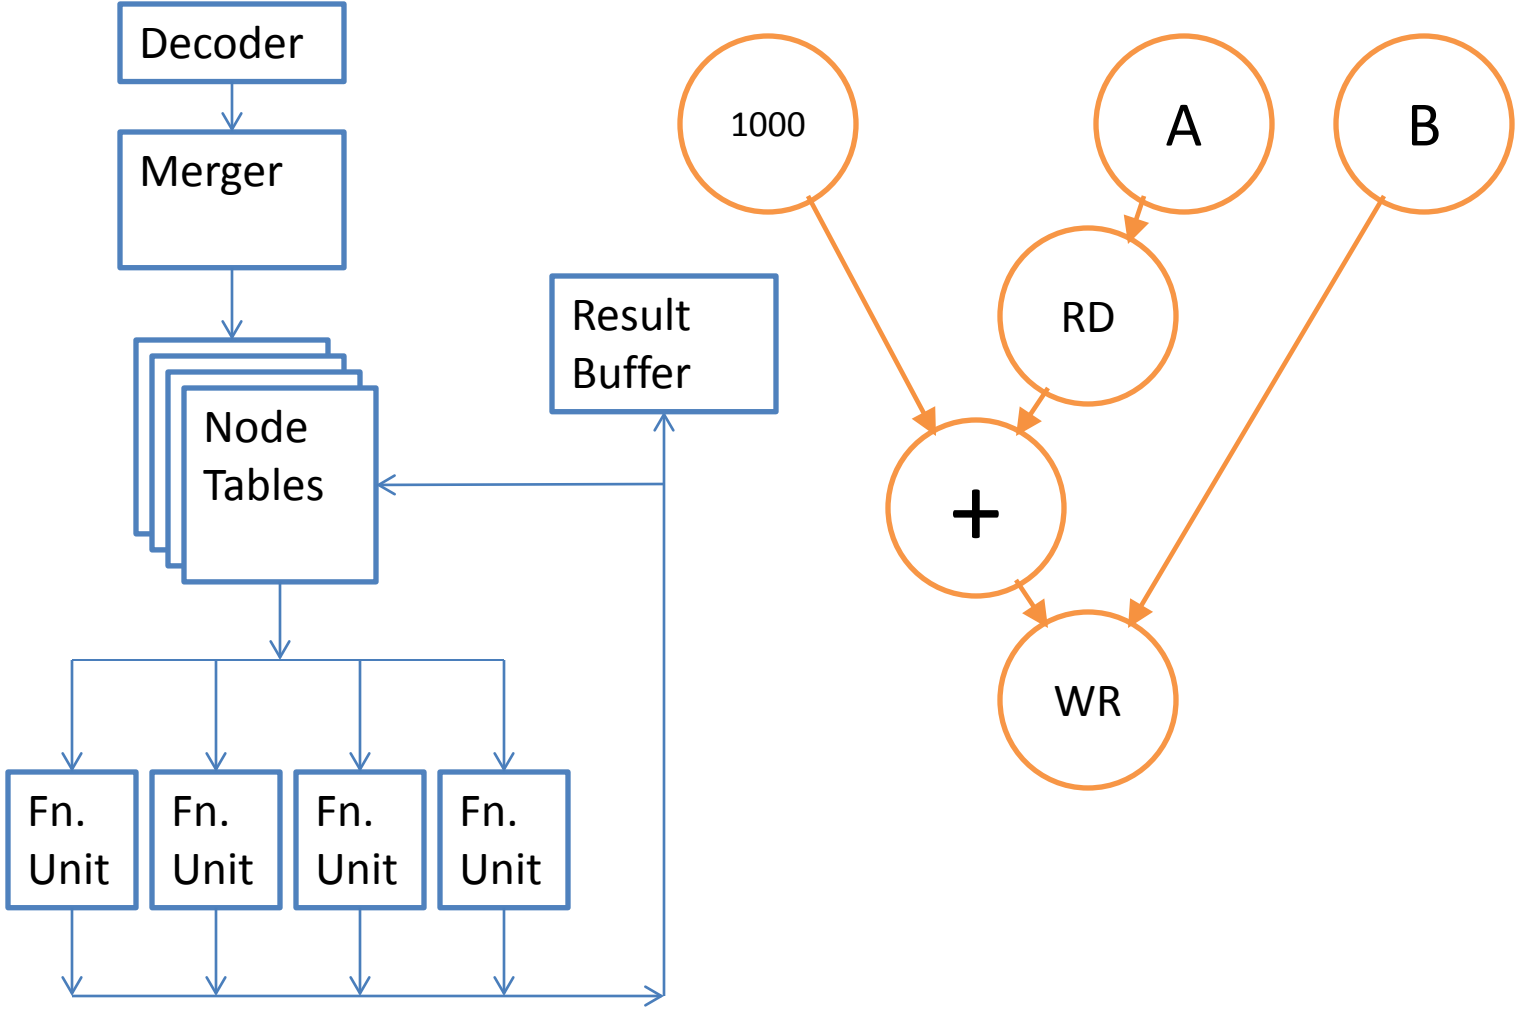
# High Performance Substrate

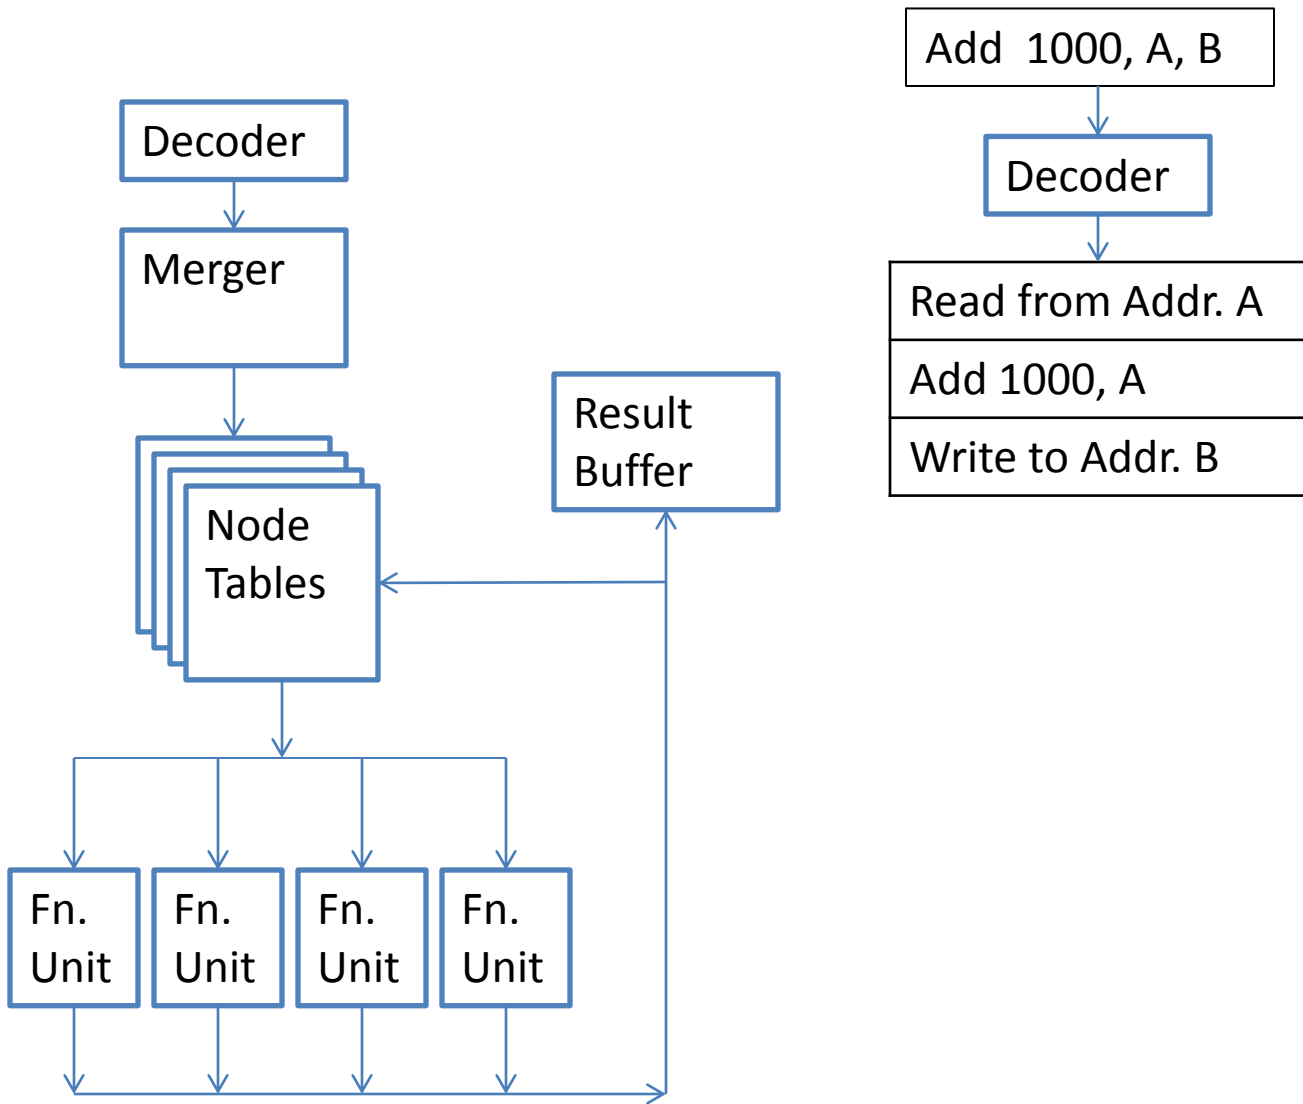
- ~20 years after Tomasulo's seminal paper, slightly different game.
- Three Tiers to Optimize
  - Global Parallelism
  - Sequential Flow
  - Local Parallelism
- Exploit local parallelism with a very small 'window' of instructions at the microarchitecture level

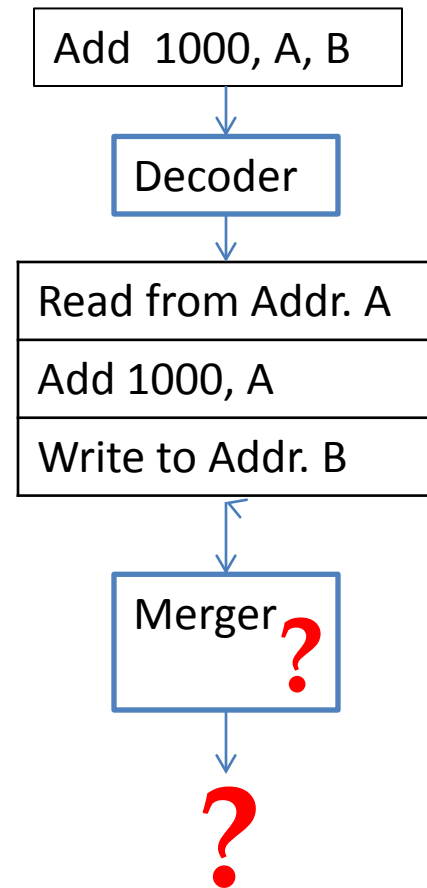
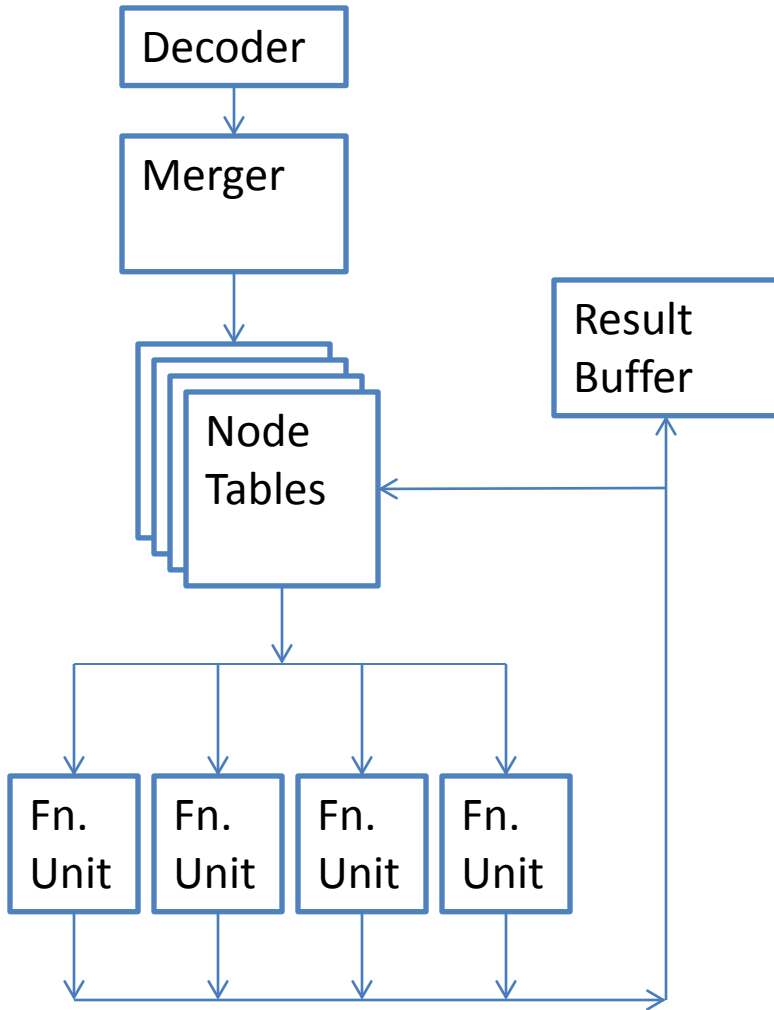






Add 1000, A, B



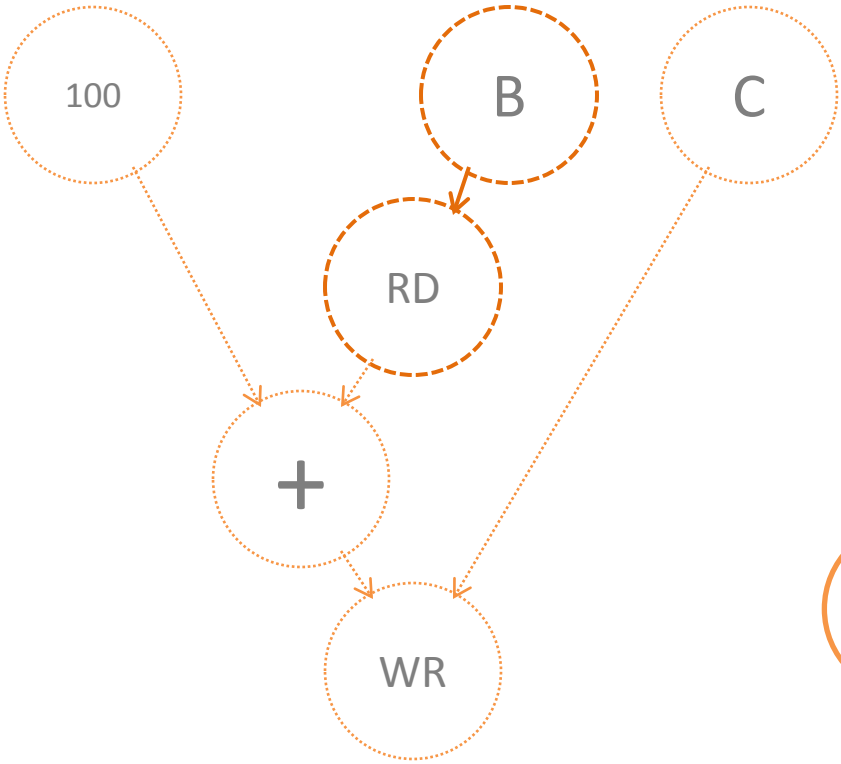




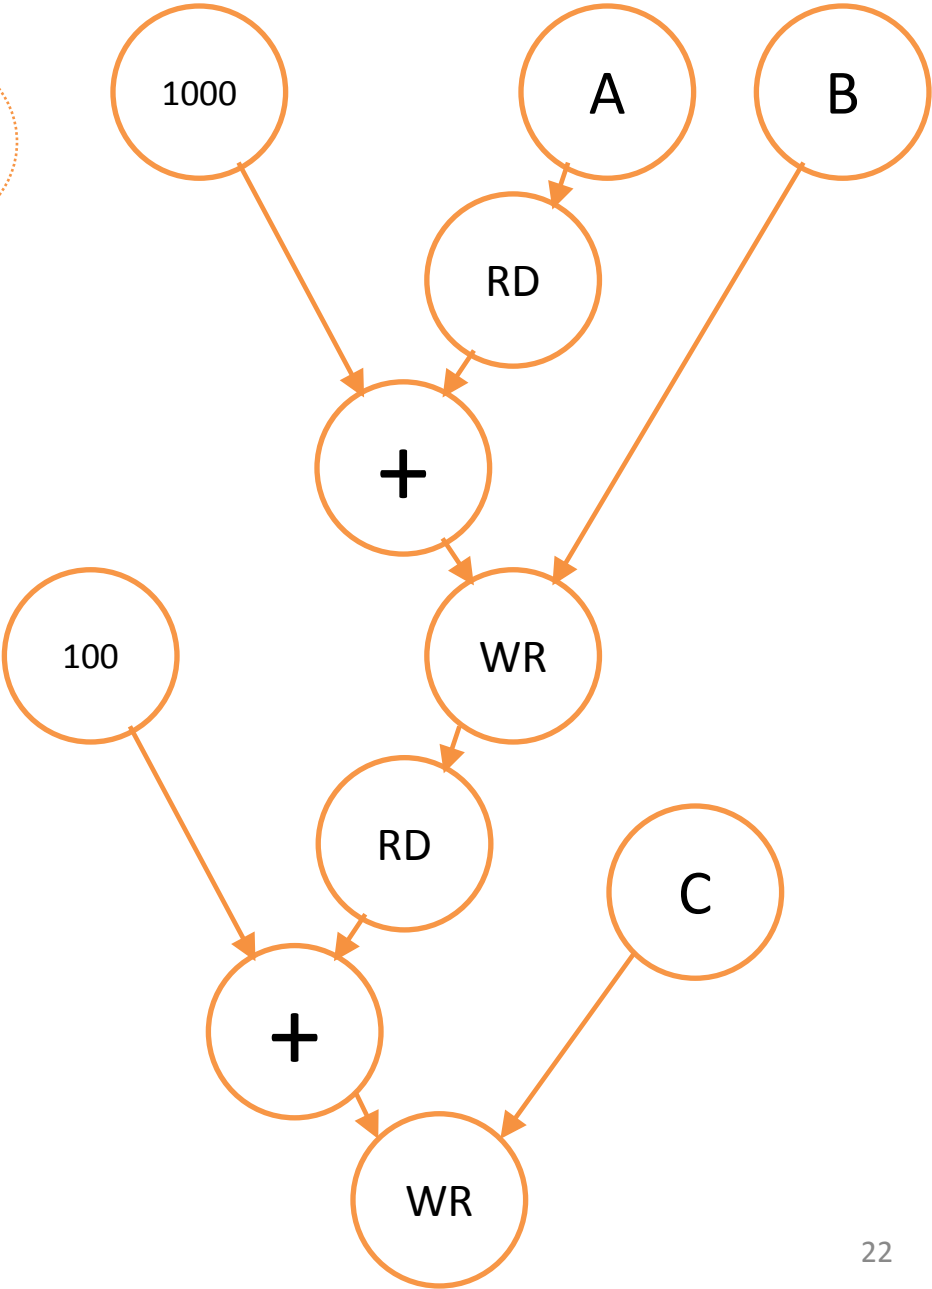
Add 1000, A, B
Add 100, B, C





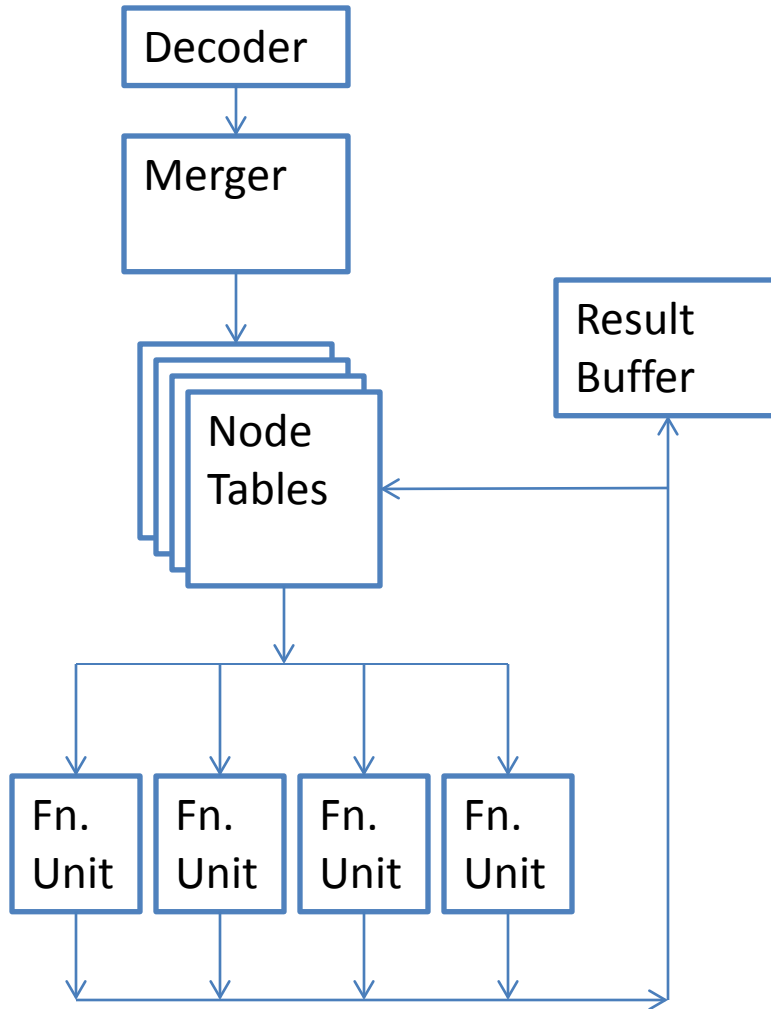
Add 100, B, C



Add 1000, A, B



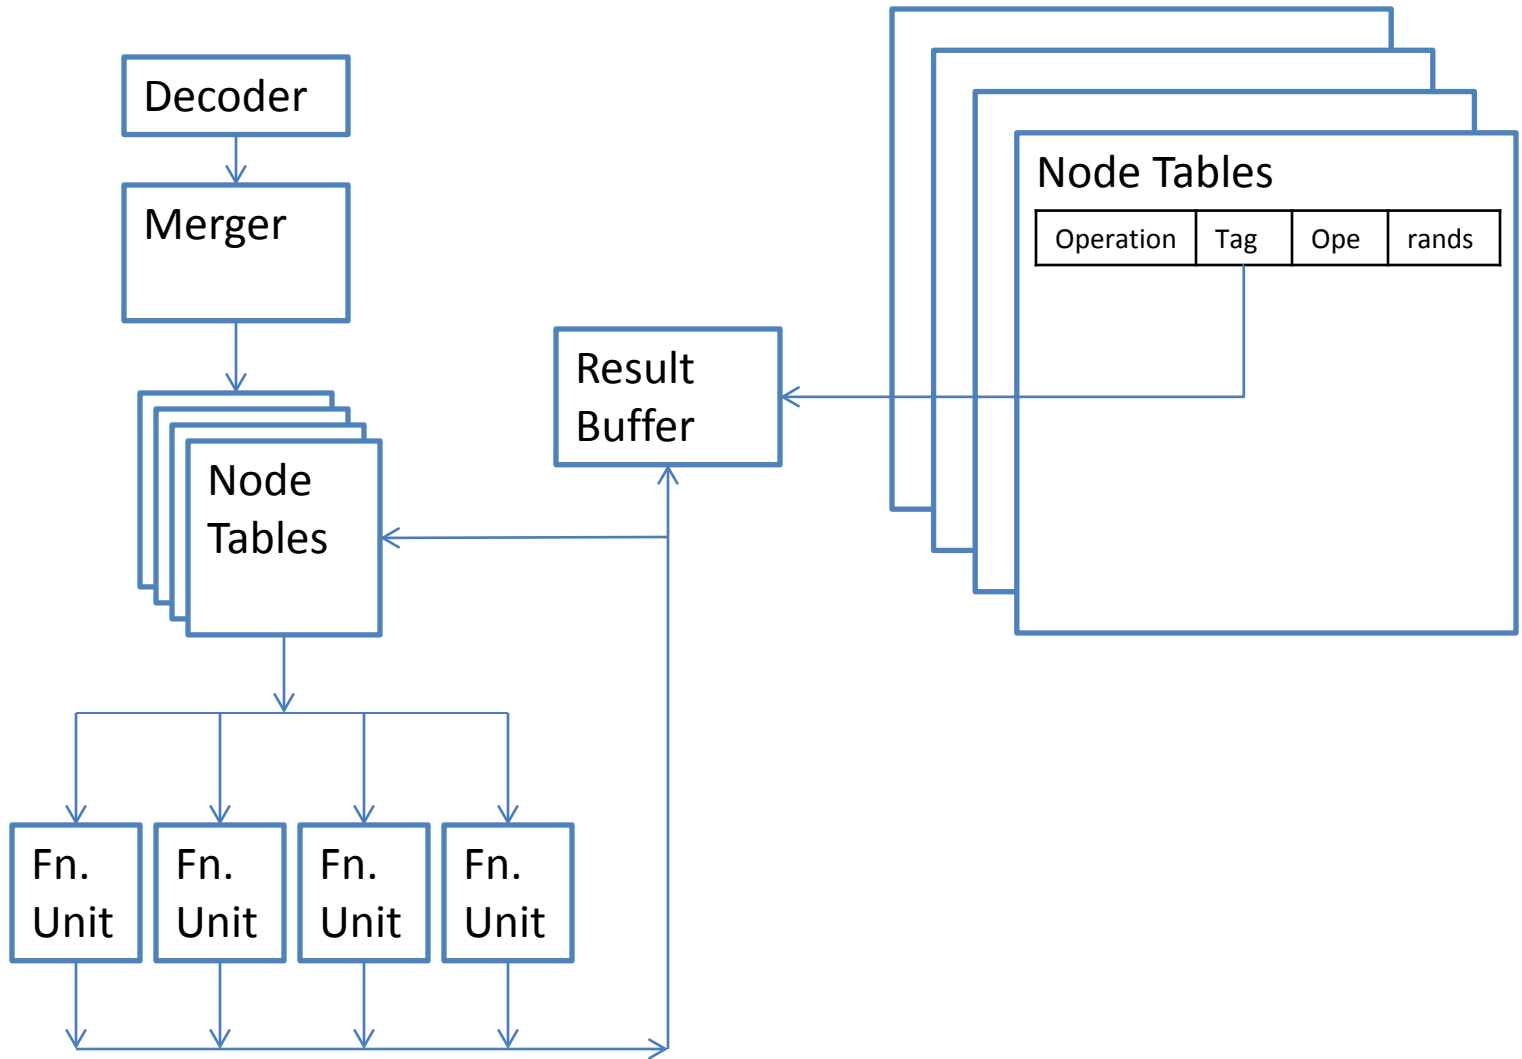


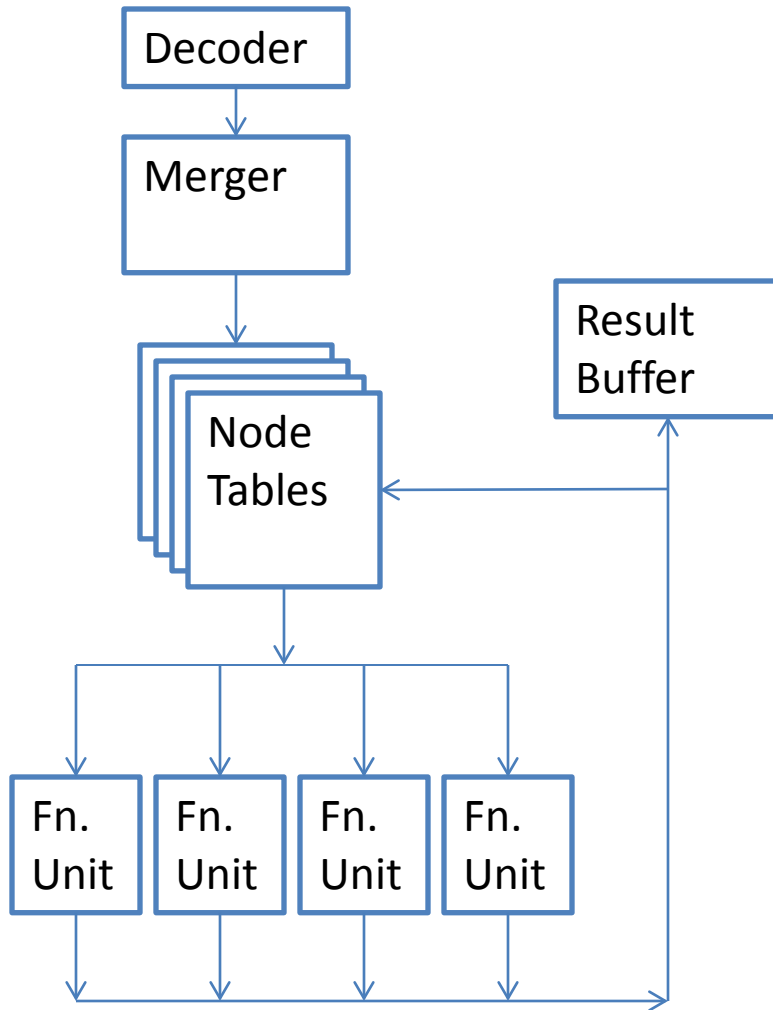


## Still Unknown:

- Format of the Node Table
- Finding Dependencies
- Scheduling Nodes

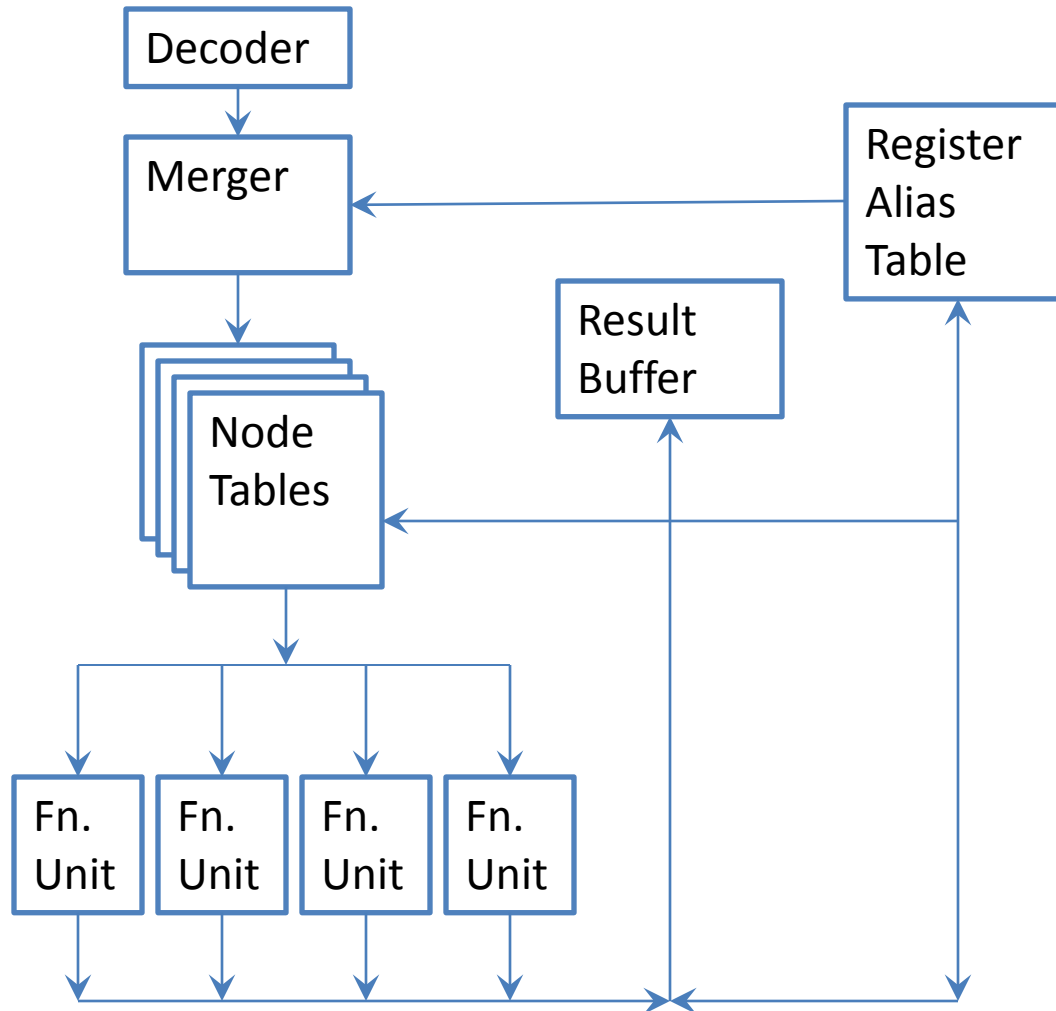




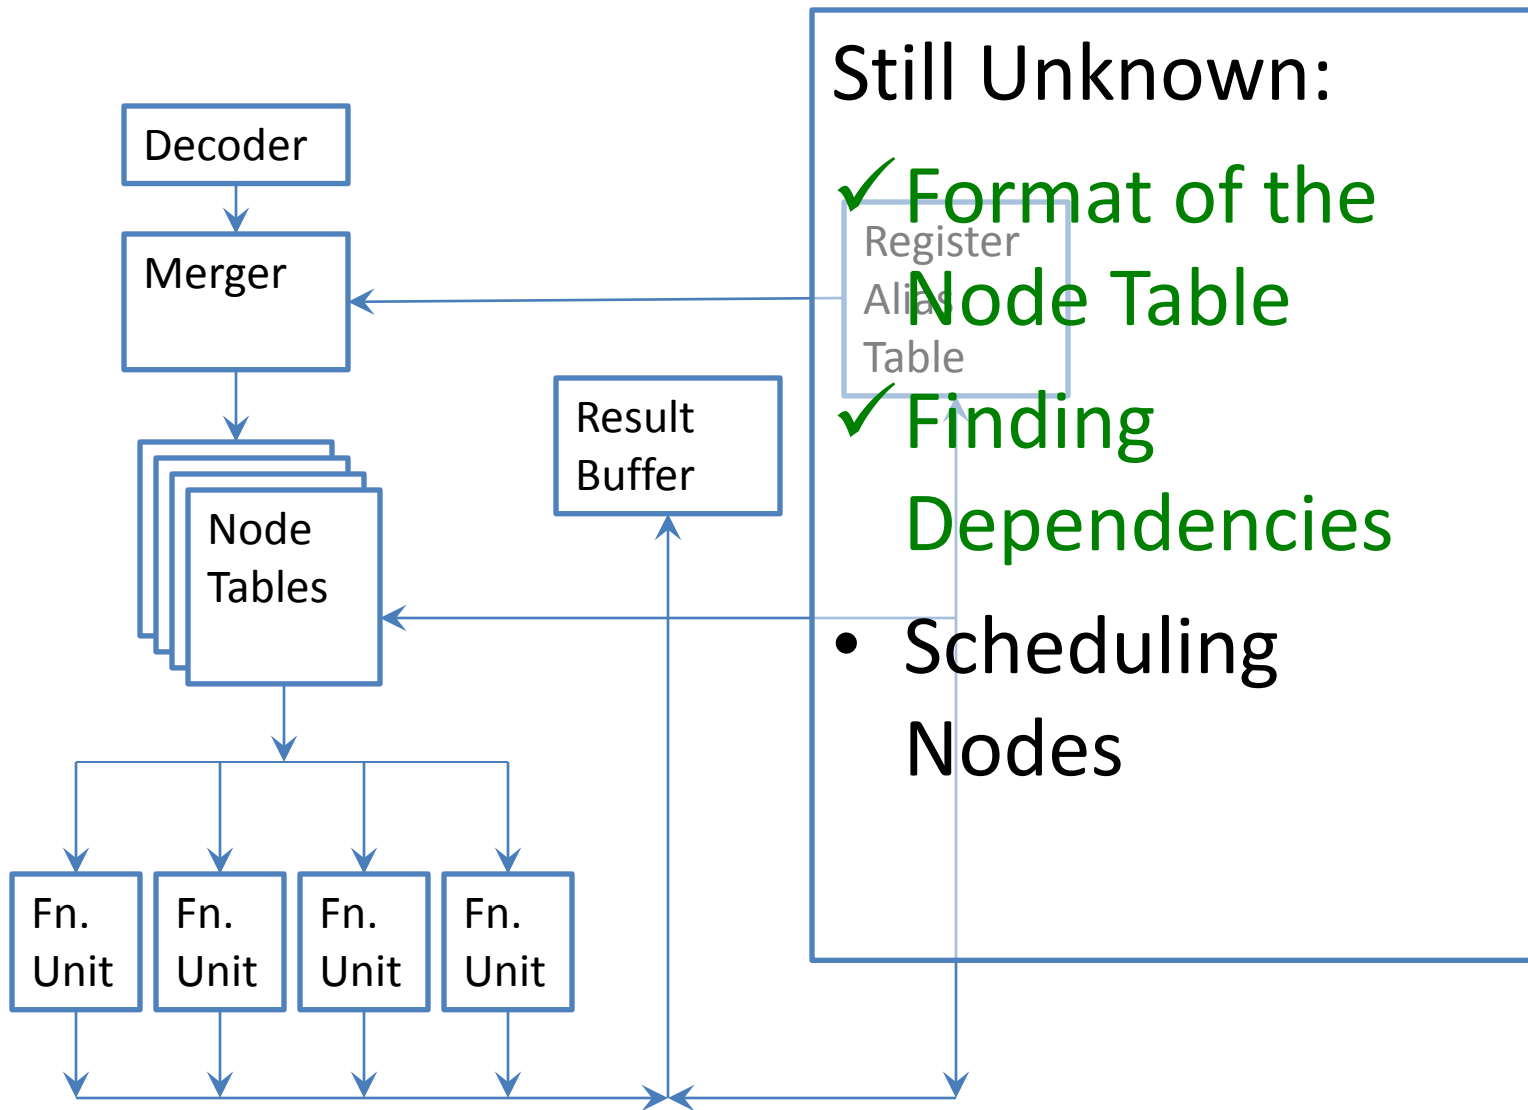


## Still Unknown:

- ✓ Format of the Node Table
- Finding Dependencies
- Scheduling Nodes



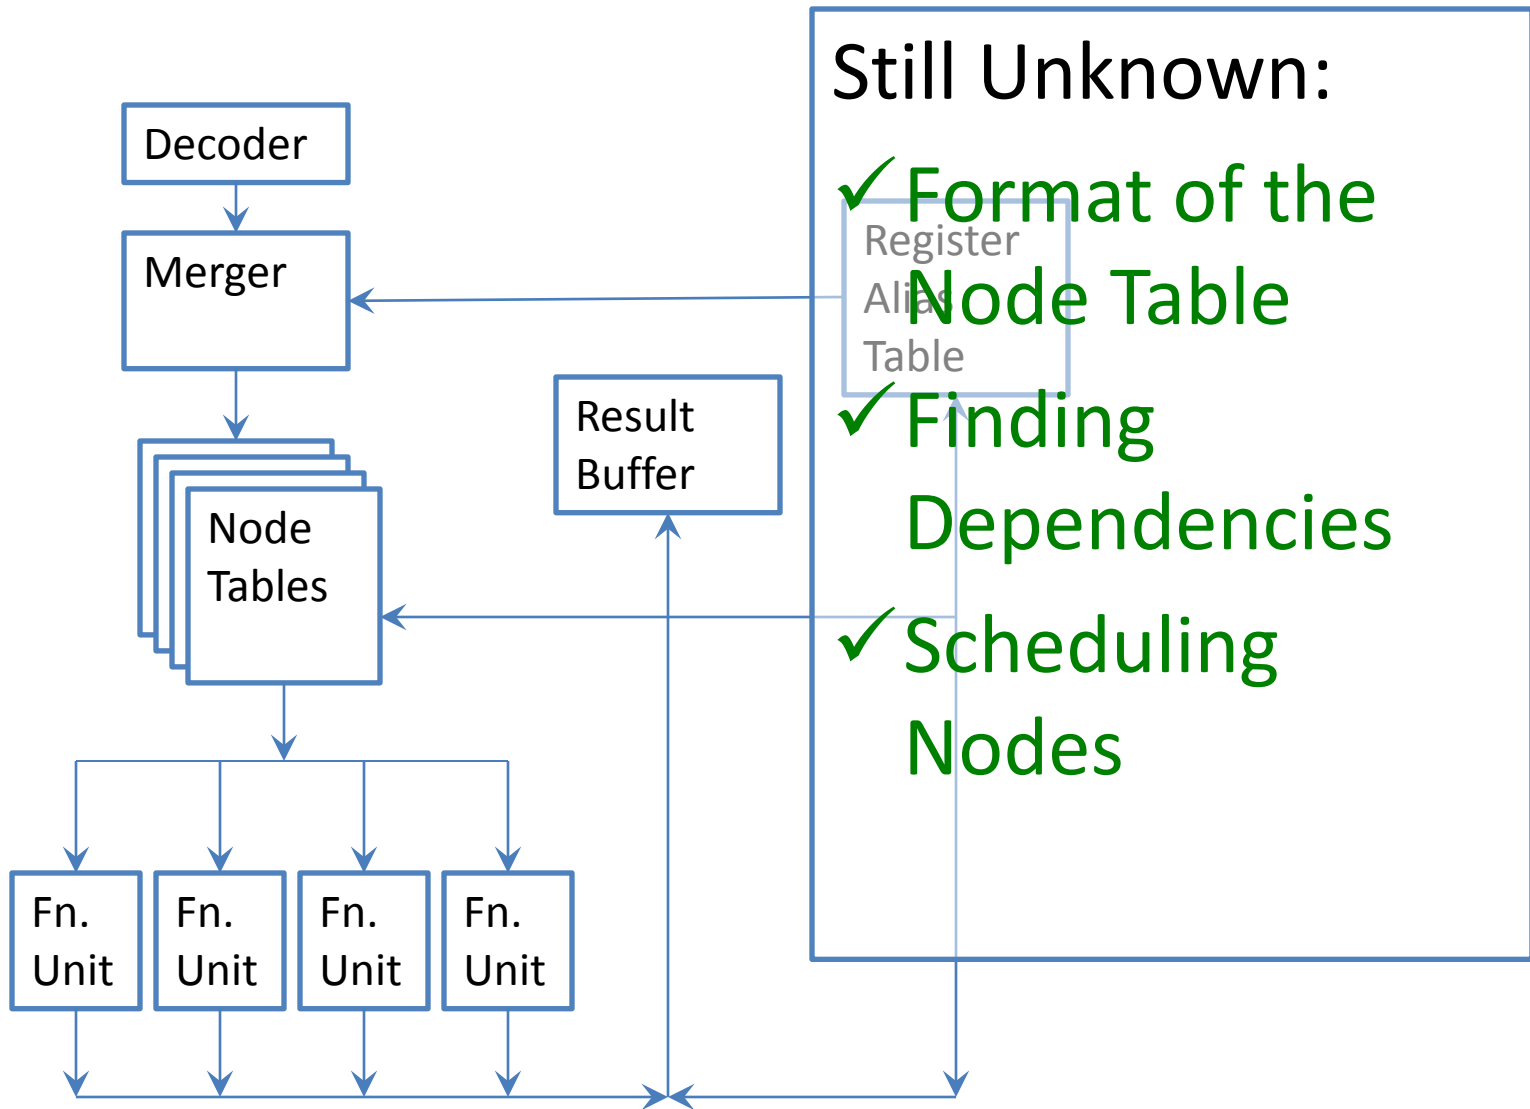
What about Memory Access?

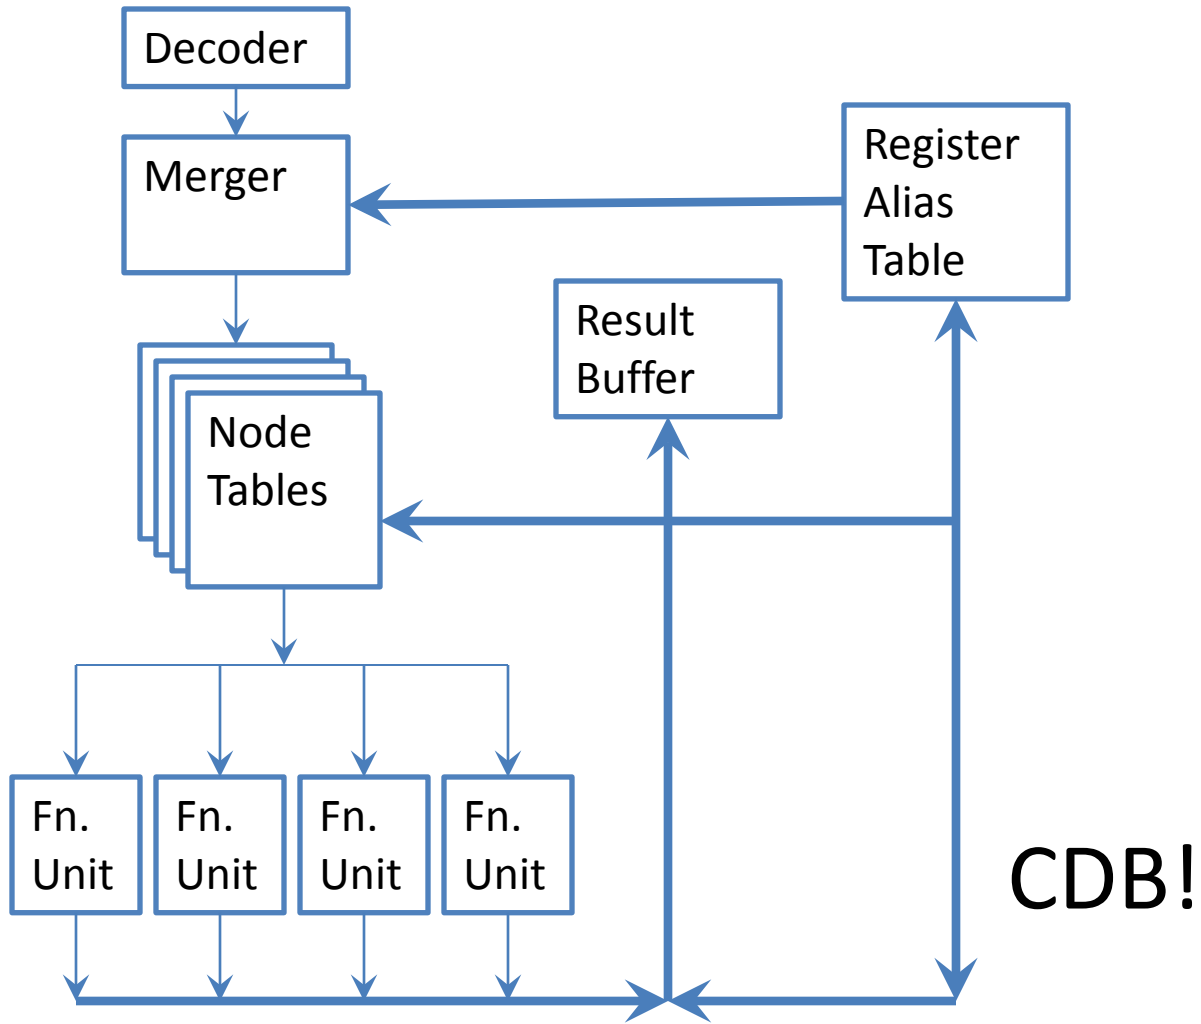


## Scheduling

- Node is 'ready to fire' when Ready Bits of all operands are set.
- Oldest 'Fires' when a Functional Unit is ready.

\* Can the scheduling make smarter choices?





# Advantages over Tomasulo's Algorithm

- No 'renaming' involved, register alias table.
  - Eliminates anti and output dependencies *without* messy renaming schemes.
- Don't need to queue instructions to 'reservation stations' before both source and sink are ready.
- Node tables allow an 'active window' worth of possible parallelism.



# HPSm[inimal]

- Implementation of the HPS model.
- Minimal, because of practical issues HPS did not address:
  - Branch Prediction.
  - Memory dependencies.
  - Number of nodes per instruction.

# HPSm[minimal]

- Implementation of the HPS model.
- Minimal, because of practical issues HPS did not address:
  - Branch Prediction.
    - Fixed to 1 unresolved prediction at a time.
  - Memory dependencies.
    - Fire oldest writes, then oldest reads.
  - Number of nodes per instruction.
    - At most two.

# Wrong Predictions and Exceptions

- We are executing out of order.
  - *What happens when the executed instructions shouldn't have been?*
  - *What happens if an exception is thrown?*
- Solution: Register Alias Table has backups

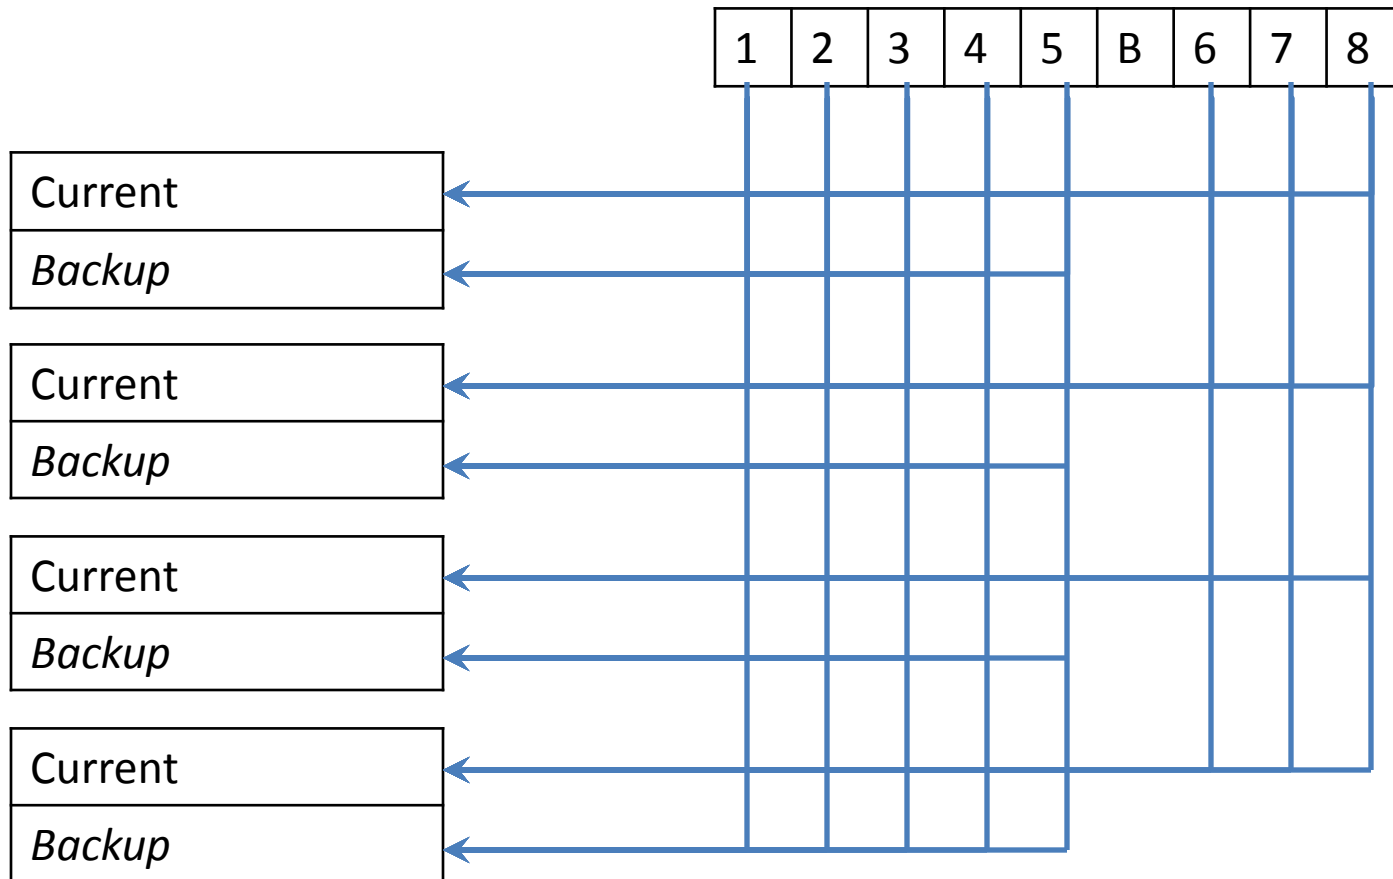
Current
<i>Backup</i>

Current
<i>Backup</i>

1	2	3	4	5	B	6	7	8
---	---	---	---	---	---	---	---	---

# Wrong Predictions and Exceptions

When the instructions are decoded:



# Wrong Predictions and Exceptions

When they are executed (out of order):

1	2	3	4	5	B	6	7	8
---	---	---	---	---	---	---	---	---

Current
<i>Backup</i>

Current
<i>Backup</i>

Current
<i>Backup</i>

Current
<i>Backup</i>

# Wrong Predictions and Exceptions

When they are executed (out of order):

1	2	3	4	5	B	6	7	8
---	---	---	---	---	---	---	---	---

<b>Current</b>
<i>Backup</i>

<b>Current</b>
<i>Backup</i>

<b>Current</b>
<i>Backup</i>

<b>Current</b>
<i>Backup</i>

# Wrong Predictions and Exceptions

When they are executed (out of order):

1	2	3	4	5	B	6	7	8
---	---	---	---	---	---	---	---	---

Current
<i>Backup</i>

Current
<i>Backup</i>





Current
<i>Backup</i>

Current
<i>Backup</i>

# Wrong Predictions and Exceptions

When they are executed (out of order):

1	2	3	4	5	<b>B</b>	X	X	X
---	---	---	---	---	----------	---	---	---






# Memory Dependencies

- Algorithm:
  - Fire the oldest fire-able memory write. If none:
  - Fire the oldest fire-able memory read.
- All access addresses translated, sit in the Write Buffer.
- *Reads* check the Write Buffer before going to memory.
- Write to memory only when the instruction RETIRES.

# HPSm Results

- Comparison against the RISC II, with both non-optimizing and optimizing compilers



# Out of Order: *Today*

- Sandy Bridge, POWER, Bulldozer and Bobcat all have **Awesome** out of order execution capabilities
  - Use Physical Register Files for renaming.
- Cortex A9, over a year ago – 1<sup>st</sup> mobile CPU with an OOO Execution Engine.
- Superscalar, OOO, Register Renaming – We're hitting an 'ILP Wall'

**Thank you**

**Any Questions?**