

CS3410

Guest Lecture

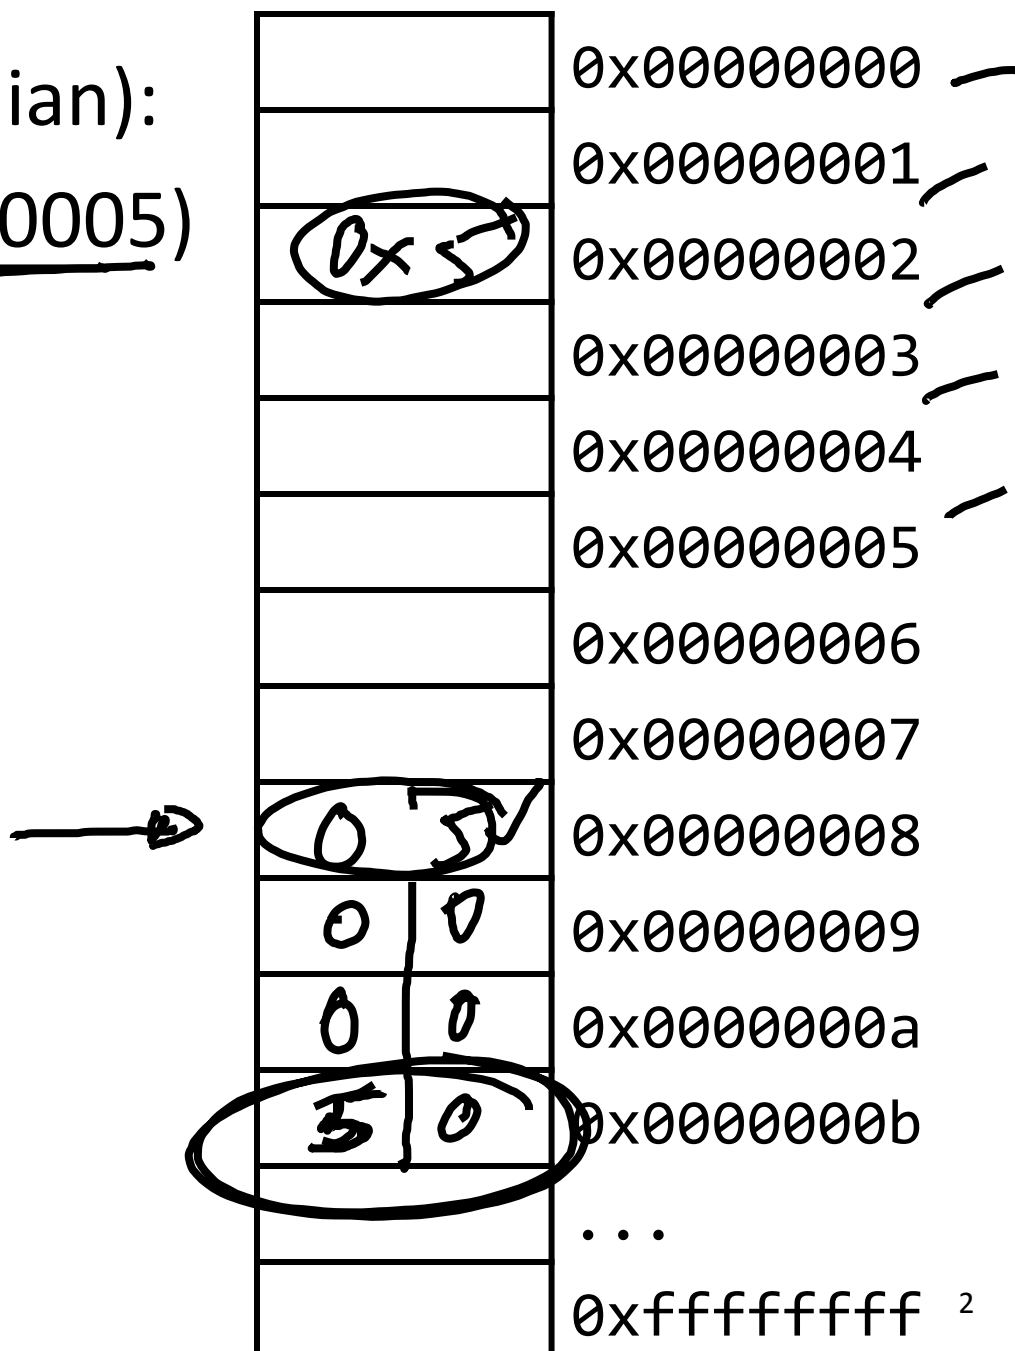
A Simple CPU: remaining branch instructions  
CPU Performance  
Pipelined CPU

Tudor Marian


Examples (big/little endian):  
 # r5 contains 5 (0x00000005)

sb r5, 2(r0) ✓  
 lb r6, 2(r0)

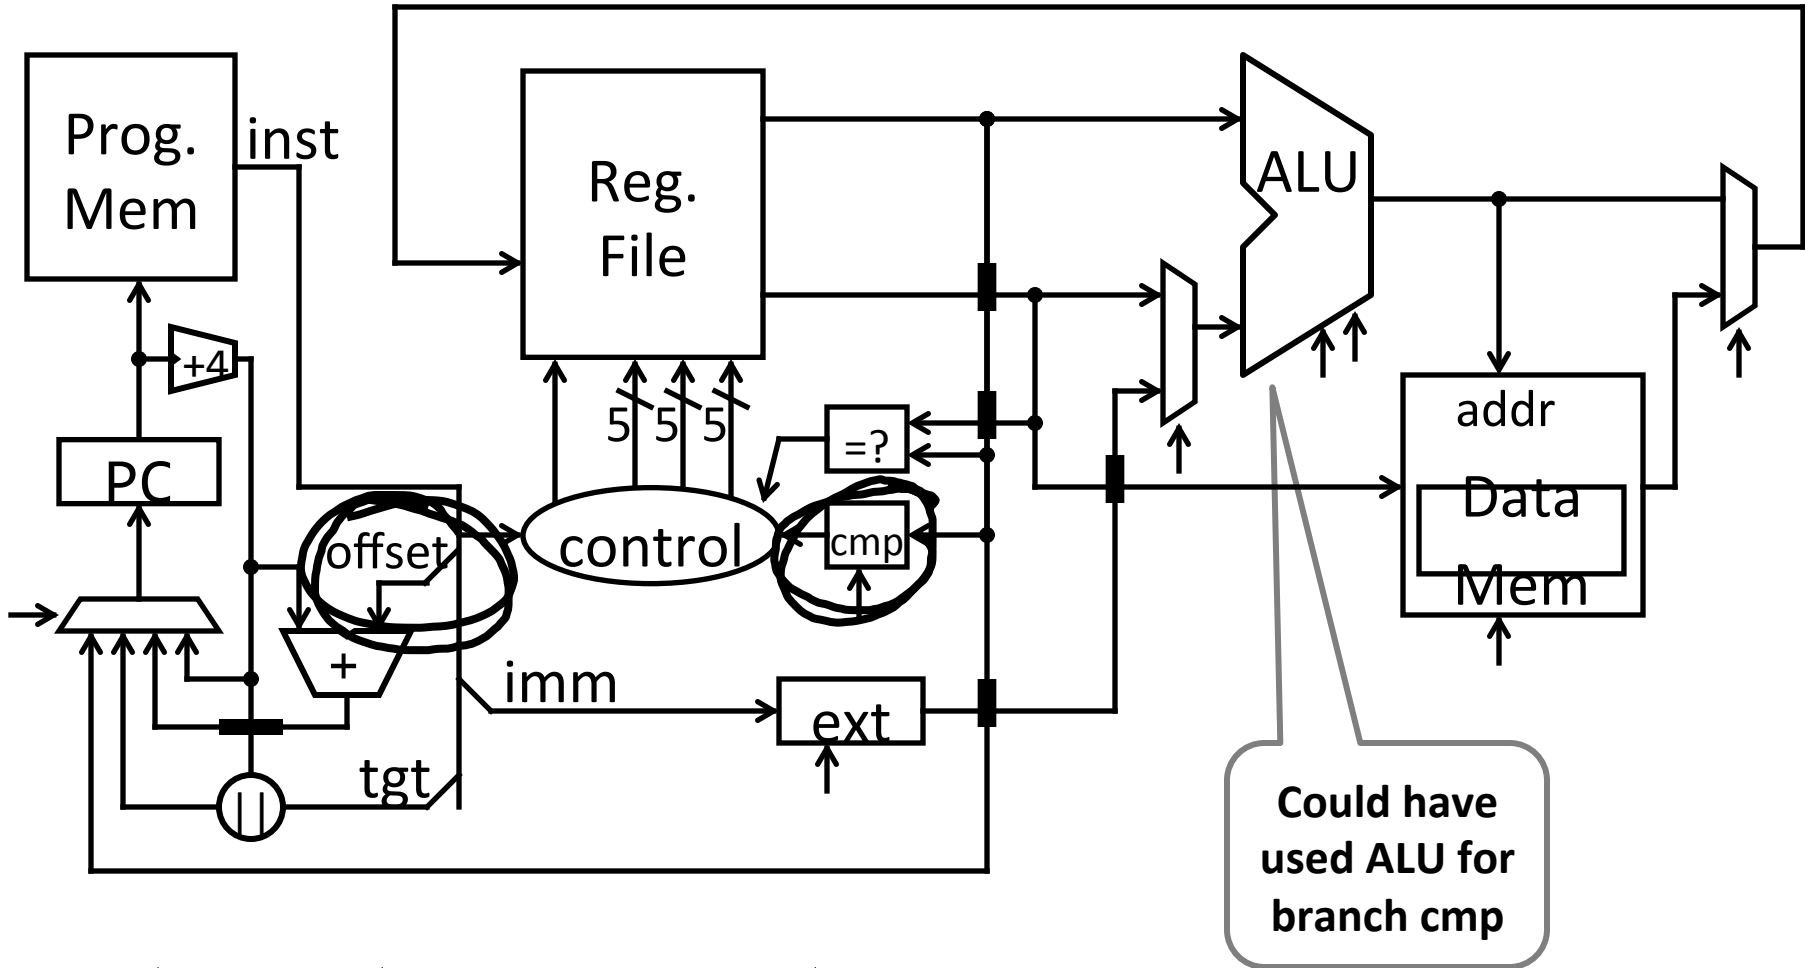
sw r5, 8(r0)  
~~lb r7, 8(r0)~~  
lb r8, 11(r0)



# Conditional Jumps (cont.)



op	subop	mnemonic	description
0x1	0x0	<u>BLTZ</u> rs, offset	if $R[rs] < 0$ then $PC = PC + 4 + (\text{offset} \ll 2)$
0x1	0x1	<u>BGEZ</u> rs, offset	if $R[rs] \geq 0$ then $PC = PC + 4 + (\text{offset} \ll 2)$
0x6	0x0	<u>BLEZ</u> rs, offset	if $R[rs] \leq 0$ then $PC = PC + 4 + (\text{offset} \ll 2)$
0x7	0x0	<u>BGTZ</u> rs, offset	if $R[rs] > 0$ then $PC = PC + 4 + (\text{offset} \ll 2)$



op	subop	mnemonic	description
0x1	0x0	BLTZ rs, offset	if $R[rs] < 0$ then $PC = PC + 4 + (\text{offset} \ll 2)$
0x1	0x1	BGEZ rs, offset	if $R[rs] \geq 0$ then $PC = PC + 4 + (\text{offset} \ll 2)$
0x1	0x2	BLTZ rs, offset	if $R[rs] < 0$ then $PC = PC + 4 + (\text{offset} \ll 2)$

# Function/procedure calls

```
00001100000001001000011000000010
```

```
|-----|-----|
```

op

immediate

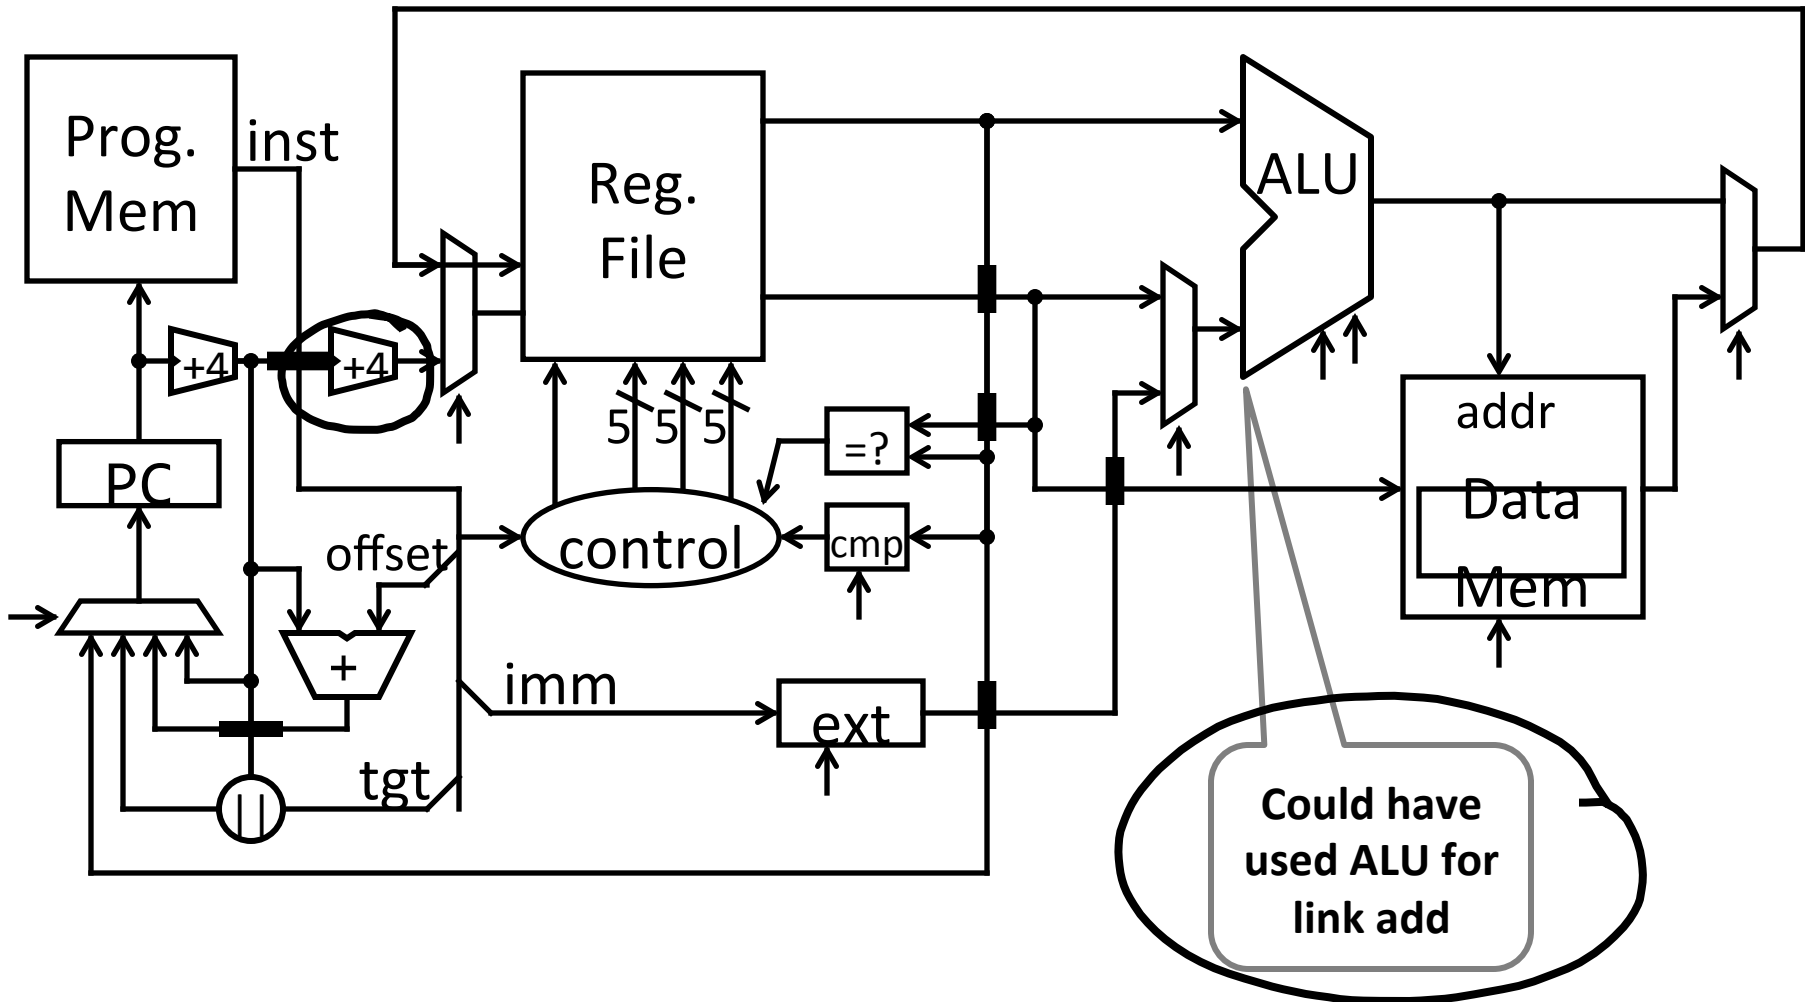
6 bits

26 bits


 J-Type

op	mnemonic	description
0x3	JAL target	r31 = PC+8 (+8 due to branch delay slot) PC = (PC+4)    (target << 2)

op	mnemonic	description
0x2	J target	PC = (PC+4)    (target << 2)



op	mnemonic	description
0x3	JAL target	$r31 = PC + 8$ (+8 due to branch delay slot) $PC = (PC + 4) \quad    \quad (target \ll 2)$

# Performance

See: P&H 1.4

# What to look for in a computer system?

- Correctness: negotiable?
- Cost
  - purchase cost =  $f(\text{silicon size} = \text{gate count, economics})$
  - operating cost =  $f(\text{energy, cooling})$
  - operating cost  $\geq$  purchase cost
- Efficiency
  - power =  $f(\text{transistor usage, voltage, wire size, clock rate, ...})$
  - heat =  $f(\text{power})$
  - Intel Core i7 Bloomfield: 130 Watts
  - AMD Turion: 35 Watts
  - Intel Core 2 Solo: 5.5 Watts
  - Cortex-A9 Dual Core @800MHz: 0.4 Watts
- Performance
- Other: availability, size, greenness, features, ...



# How to measure performance?

GHz ;  $\frac{1}{2}$  ; 1GHz

→. *latency*

GHz (billions of cycles per second)

MIPS (millions of instructions per second)

MFLOPS (millions of floating point operations per second)

benchmarks (SPEC, TPC, ...)

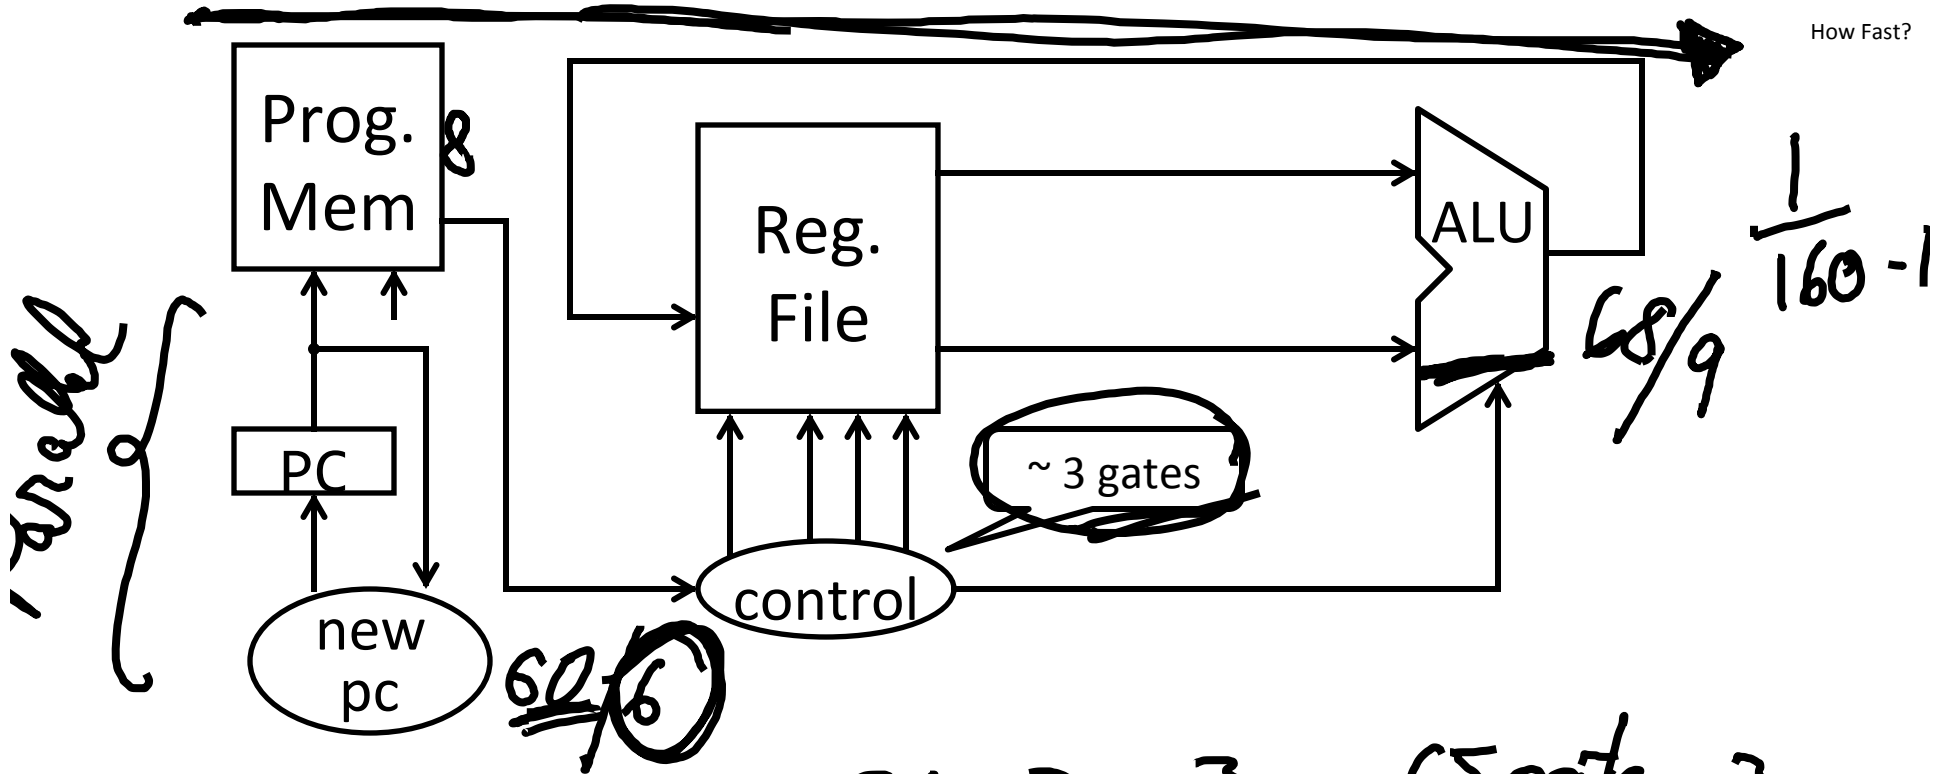
MIPS →  $10^6$  instr/sec

MFLOPS

## Metrics

latency: how long to finish my program

throughput: how much work finished per unit time



Assumptions:

- alu: 62 bit ripple carry + some muxes
- next PC: 50 bit ripple carry
- control: minimized for delay (~3 gates)
- transistors: 2 ns per gate
- prog. memory: 16 ns (as much as 8 gates)
- register file: 2 ns access
- ignore wires, register setup time

Better:

- alu: 32 bit carry lookahead + some muxes (~9 gates)
- next PC: 30 bit carry lookahead (~6 gates)

Better Still:

- next PC: cheapest adder faster than 21 gate delays

$31 \times 2 + 3 = 65 \text{ gates} + 3 \sim$   
 $(8 + 3 + 1 + 68) = 80 \text{ gate}$   
 $8 + 3 + 9 + 1 = 21 \text{ gate}$

All signals are stable

- 80 gates => clock period of at least 160 ns, max frequency ~6MHz

Better:

- 21 gates => clock period of at least 42 ns, max frequency ~24MHz

## 32 Bit Adder Design

Ripple Carry

2-Way Carry-Skip

3-Way Carry-Skip

4-Way Carry-Skip

2-Way Look-Ahead

Split Look-Ahead

Full Look-Ahead

## Space

≈ 300 gates

≈ 360 gates

≈ 500 gates

≈ 600 gates

≈ 550 gates

≈ 800 gates

≈ 1200 gates

## Time

≈ 64 gate delays

≈ 35 gate delays

≈ 22 gate delays

≈ 18 gate delays

≈ 16 gate delays

≈ 10 gate delays

≈ 5 gate delays

## Critical Path

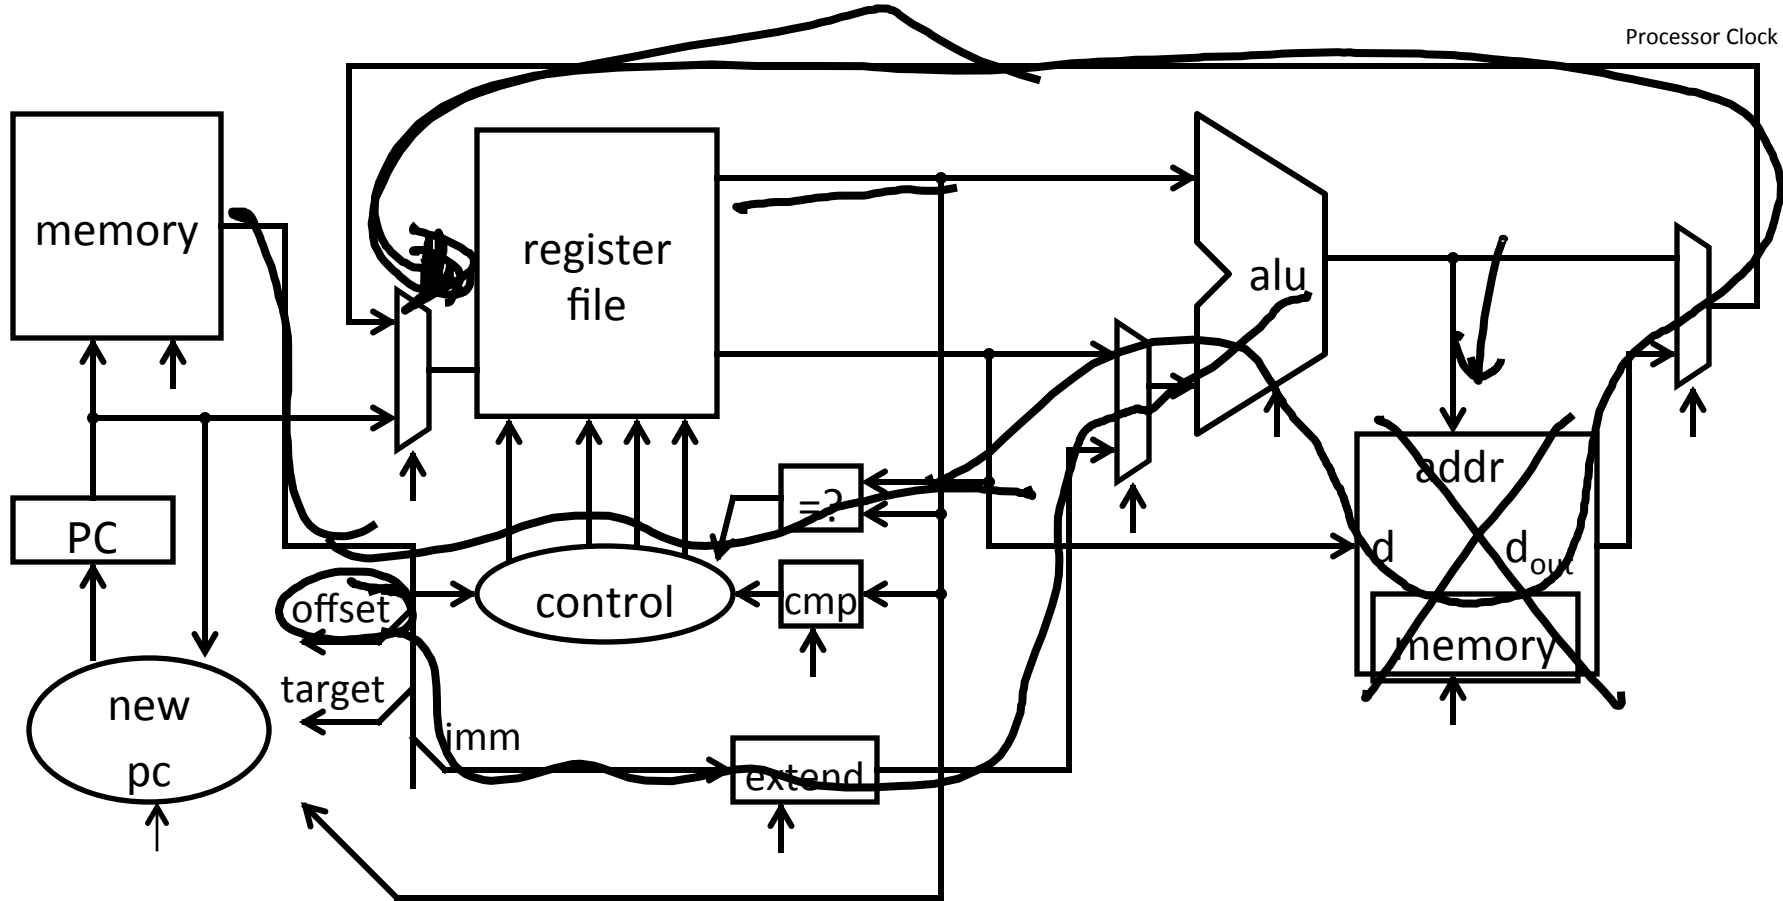
- Longest path from a register output to a register input
- Determines minimum cycle, maximum clock frequency

## Strategy 1 (we just employed)

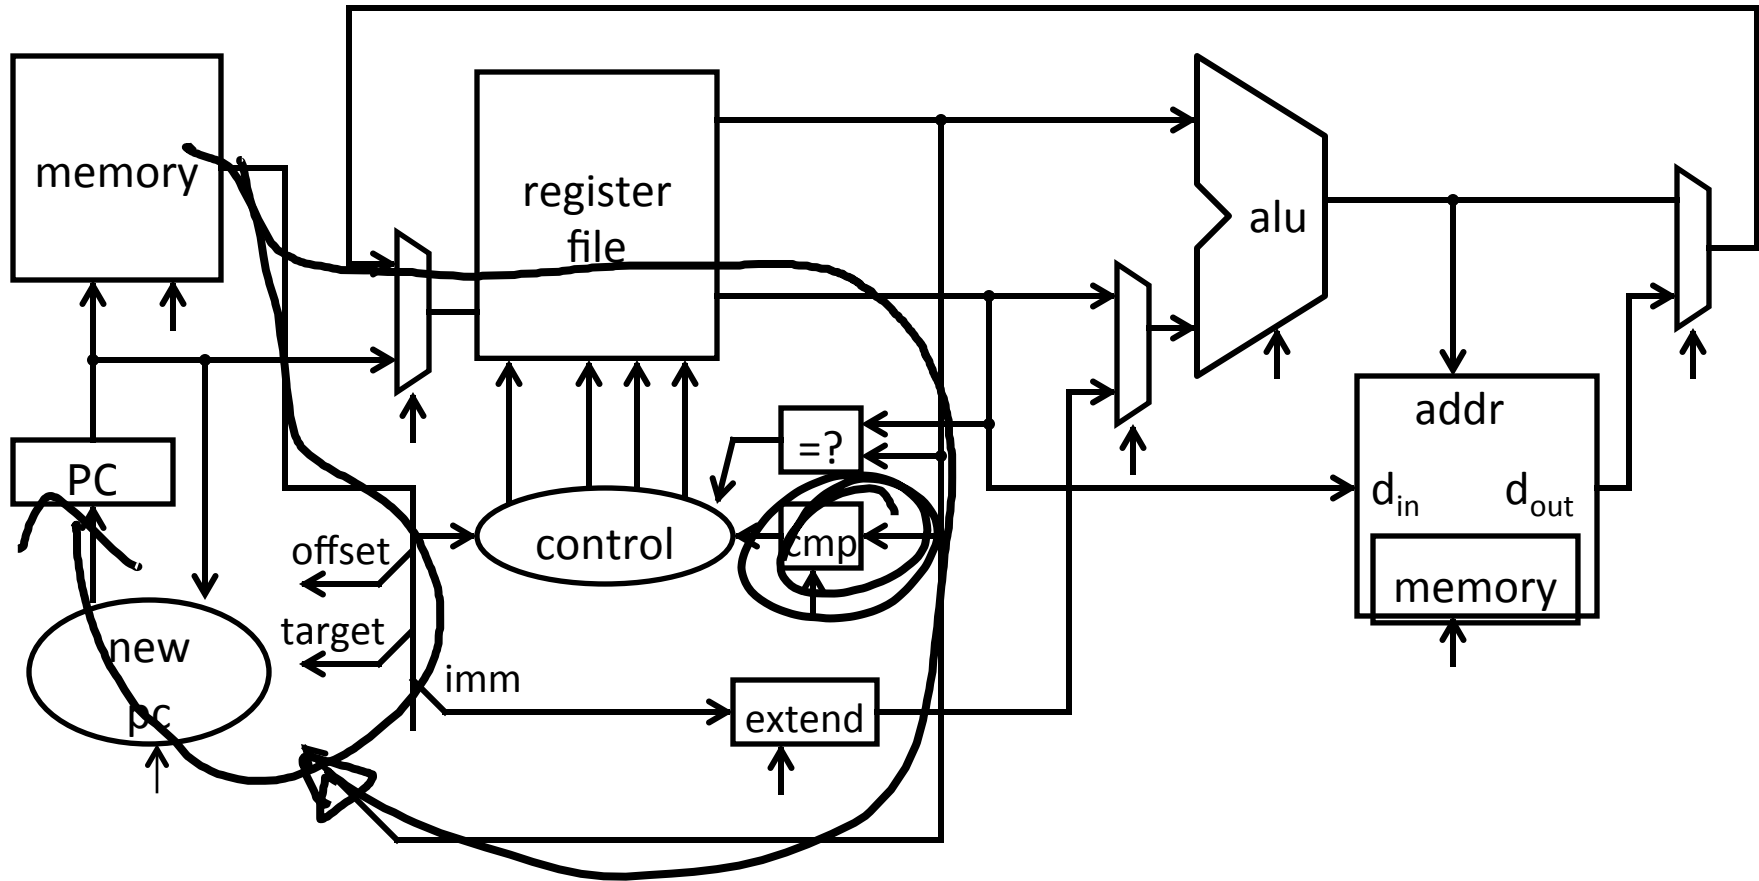
- Optimize for delay on the critical path
- Optimize for size / power / simplicity elsewhere

– next PC

Single cycle



op	mnemonic	description
0x20	LB rd, offset(rs)	$R[rd] = \text{sign\_ext}(\text{Mem}[\text{offset}+R[rs]])$
0x23	LW rd, offset(rs)	$R[rd] = \text{Mem}[\text{offset}+R[rs]]$
0x28	SB rd, offset(rs)	$\text{Mem}[\text{offset}+R[rs]] = R[rd]$
0x2h	SW rd, offset(rs)	$\text{Mem}[\text{offset}+R[rs]] = R[rd]$



<b>op</b>	<b>func</b>	<b>mnemonic</b>	<b>description</b>
0x0	0x08	JR rs	PC = R[rs]

<b>op</b>	<b>mnemonic</b>	<b>description</b>
0x2	J target	PC = (PC+4)    (target << 2)

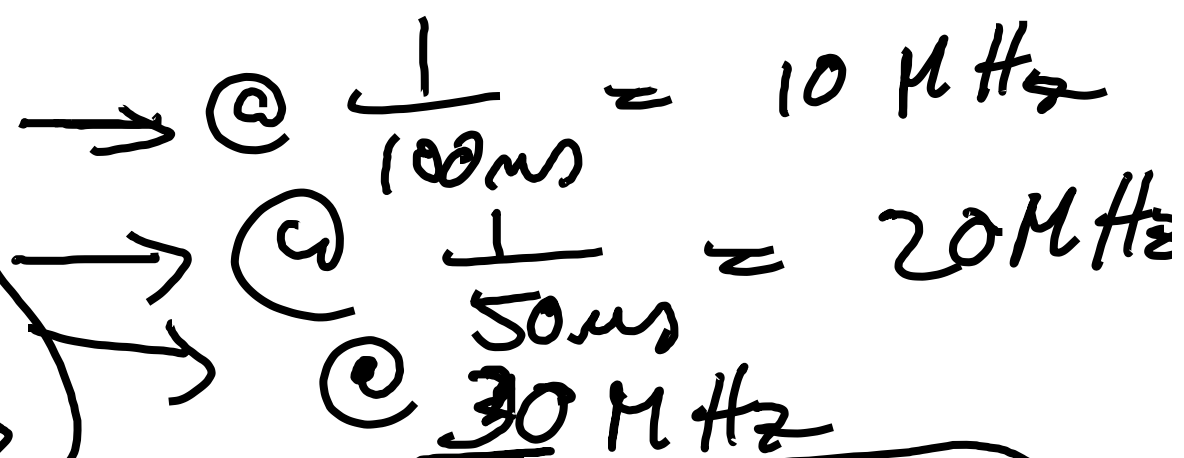
## Strategy 2

- Multiple cycles to complete a single instruction

E.g: Assume:

- load/store: 100 ns
- arithmetic: 50 ns
- branches: 33 ns

66 ns  
99 ns



Multi-Cycle CPU

- 3 cycles per load/store
- 2 cycles per arithmetic
- 1 cycle per branch

Faster than Single-Cycle CPU?  
 10 MHz (100 ns cycle) with  
 - 1 cycle per instruction

10 MHz = 10 Mi  
 1 CPU

Instruction mix for some program P, assume:

- 25% load/store ( 3 cycles / instruction)
- 60% arithmetic ( 2 cycles / instruction)
- 15% branches ( 1 cycle / instruction)

Multi-Cycle performance for program P:  $CPI = 0.25 \times 3 + \dots$

$$3 * .25 + 2 * .60 + 1 * .15 = 2.1$$

average cycles per instruction (CPI) = 2.1

Multi-Cycle @ 30 MHz  
 Single-Cycle @ 10 MHz  
 Single-Cycle @ 15 MHz

800 MHz PIII "faster" than 1 GHz P4

$$2.1 CPI \Rightarrow$$

10 MIPS  
 15 MIPS

$$\frac{30 \cdot 10^6 \text{ cy}}{2.1 \text{ cycles}} \approx 15 \text{ MIPS}$$



Goal: Make Multi-Cycle @ 30 MHz CPU (15MIPS) run 2x faster by making arithmetic instructions faster,

Instruction mix (for P):

- 25% load/store, CPI = 3
- 60% arithmetic, CPI = 2.1
- 15% branches, CPI = 1

0.25	0.75	0.75	0.75
0.60	1.2	0.6	0.15
0.15	0.15	0.15	0.15
<hr/>		<hr/>	
CPI = 2.1	1.50	1.0	

15 MIPS

$30 \text{ MHz} = 20 \text{ MIPS}$   
 $1 \text{ cycle} \dots 1.5 \text{ CPI}$   
 $30 \times 4 \text{ MHz}$

$\frac{30 \text{ MHz}}{1.05} = 30 \text{ MIPS}$

# Amdahl's Law

Execution time after improvement =

$$\frac{\text{execution time affected by improvement}}{\text{amount of improvement}} + \text{execution time unaffected}$$

*lw.*

Or:

Speedup is limited by popularity of improved feature

Corollary:

Make the common case fast


Caveat:

Law of diminishing returns


# Pipelining

See: P&H Chapter 4.5

Alice



Bob



They don't always get along...






Saw



Drill





Glue



Paint

N pieces, each built following same sequence:





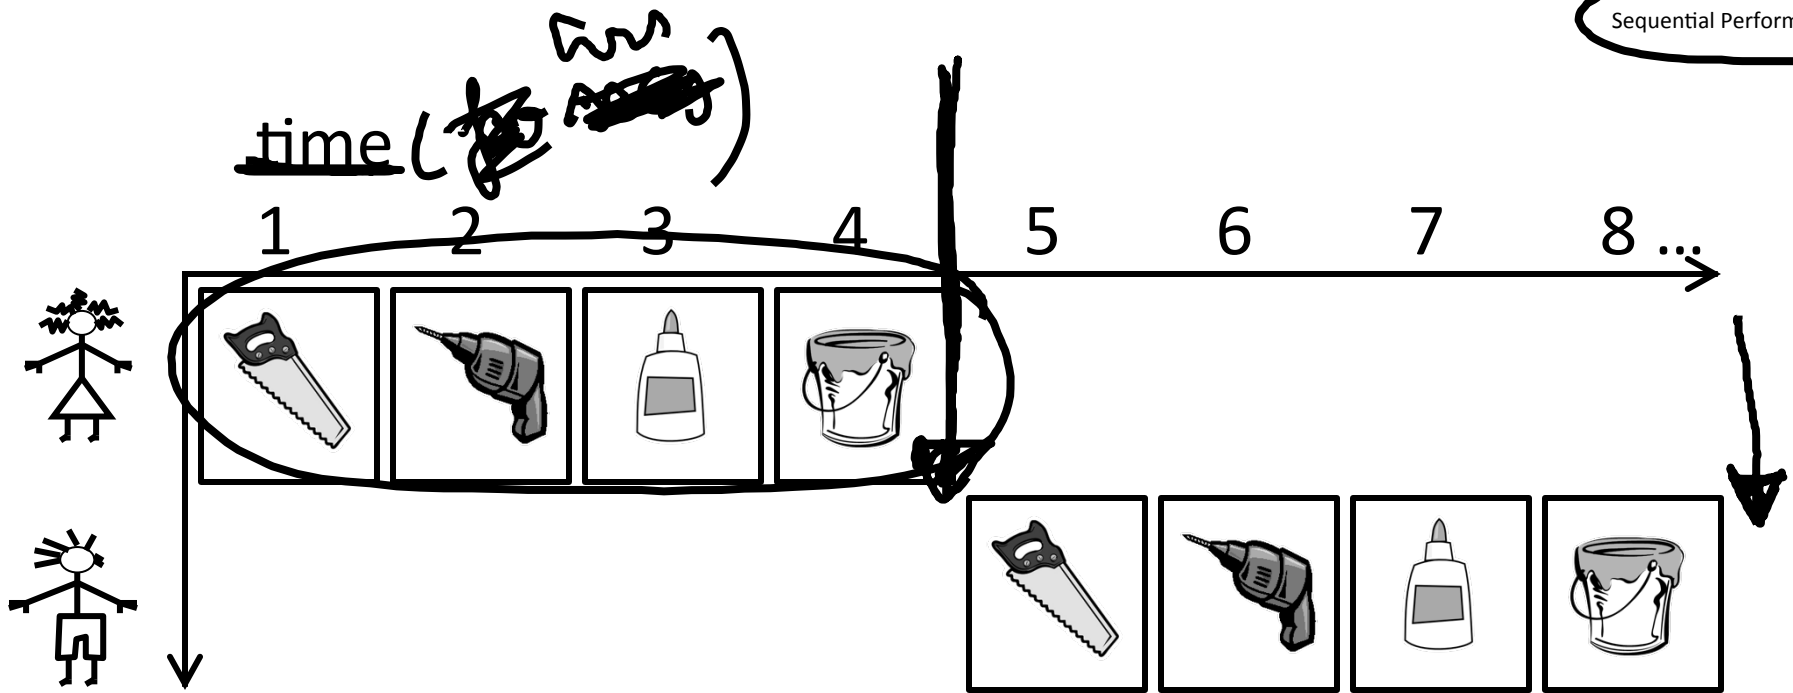
Alice owns the room

Bob can enter when Alice is finished

Repeat for remaining tasks

No possibility for conflicts






Latency: 4 hrs / task  
Throughput: 1 task / 4 hrs  
Concurrency: 1

Can we do better?

## Partition room into *stages* of a *pipeline*




One person owns a stage at a time


4 stages

4 people working simultaneously

Everyone moves right in lockstep




> Latency: 4 hrs / task  
Throughput: 1 task/hr  
Concurrency: 4



Latency:  $\infty$

Throughput:  $\frac{1 \text{ task}}{90 \text{ min}}$

Concurrency:  $\infty$



Latency:


Throughput:

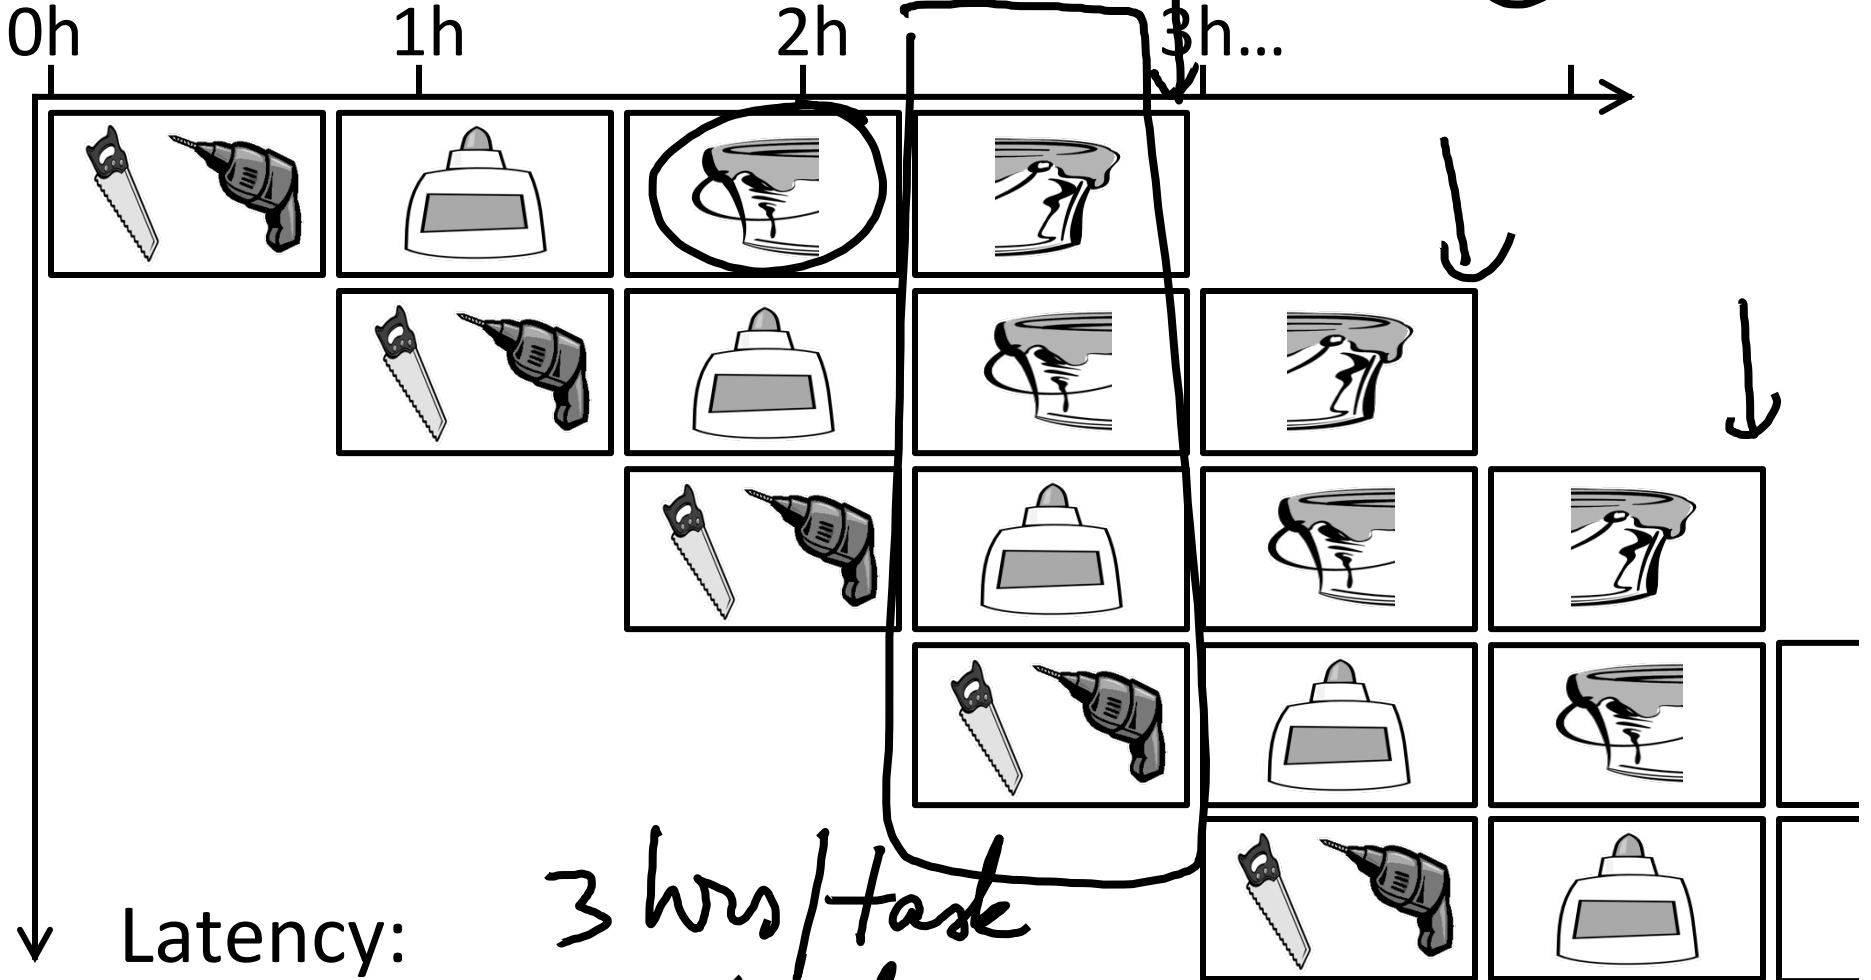
Concurrency:

$$4 \times 90 \text{ min} = 6 \text{ hrs / task}$$

$$1 \text{ task / 90}$$

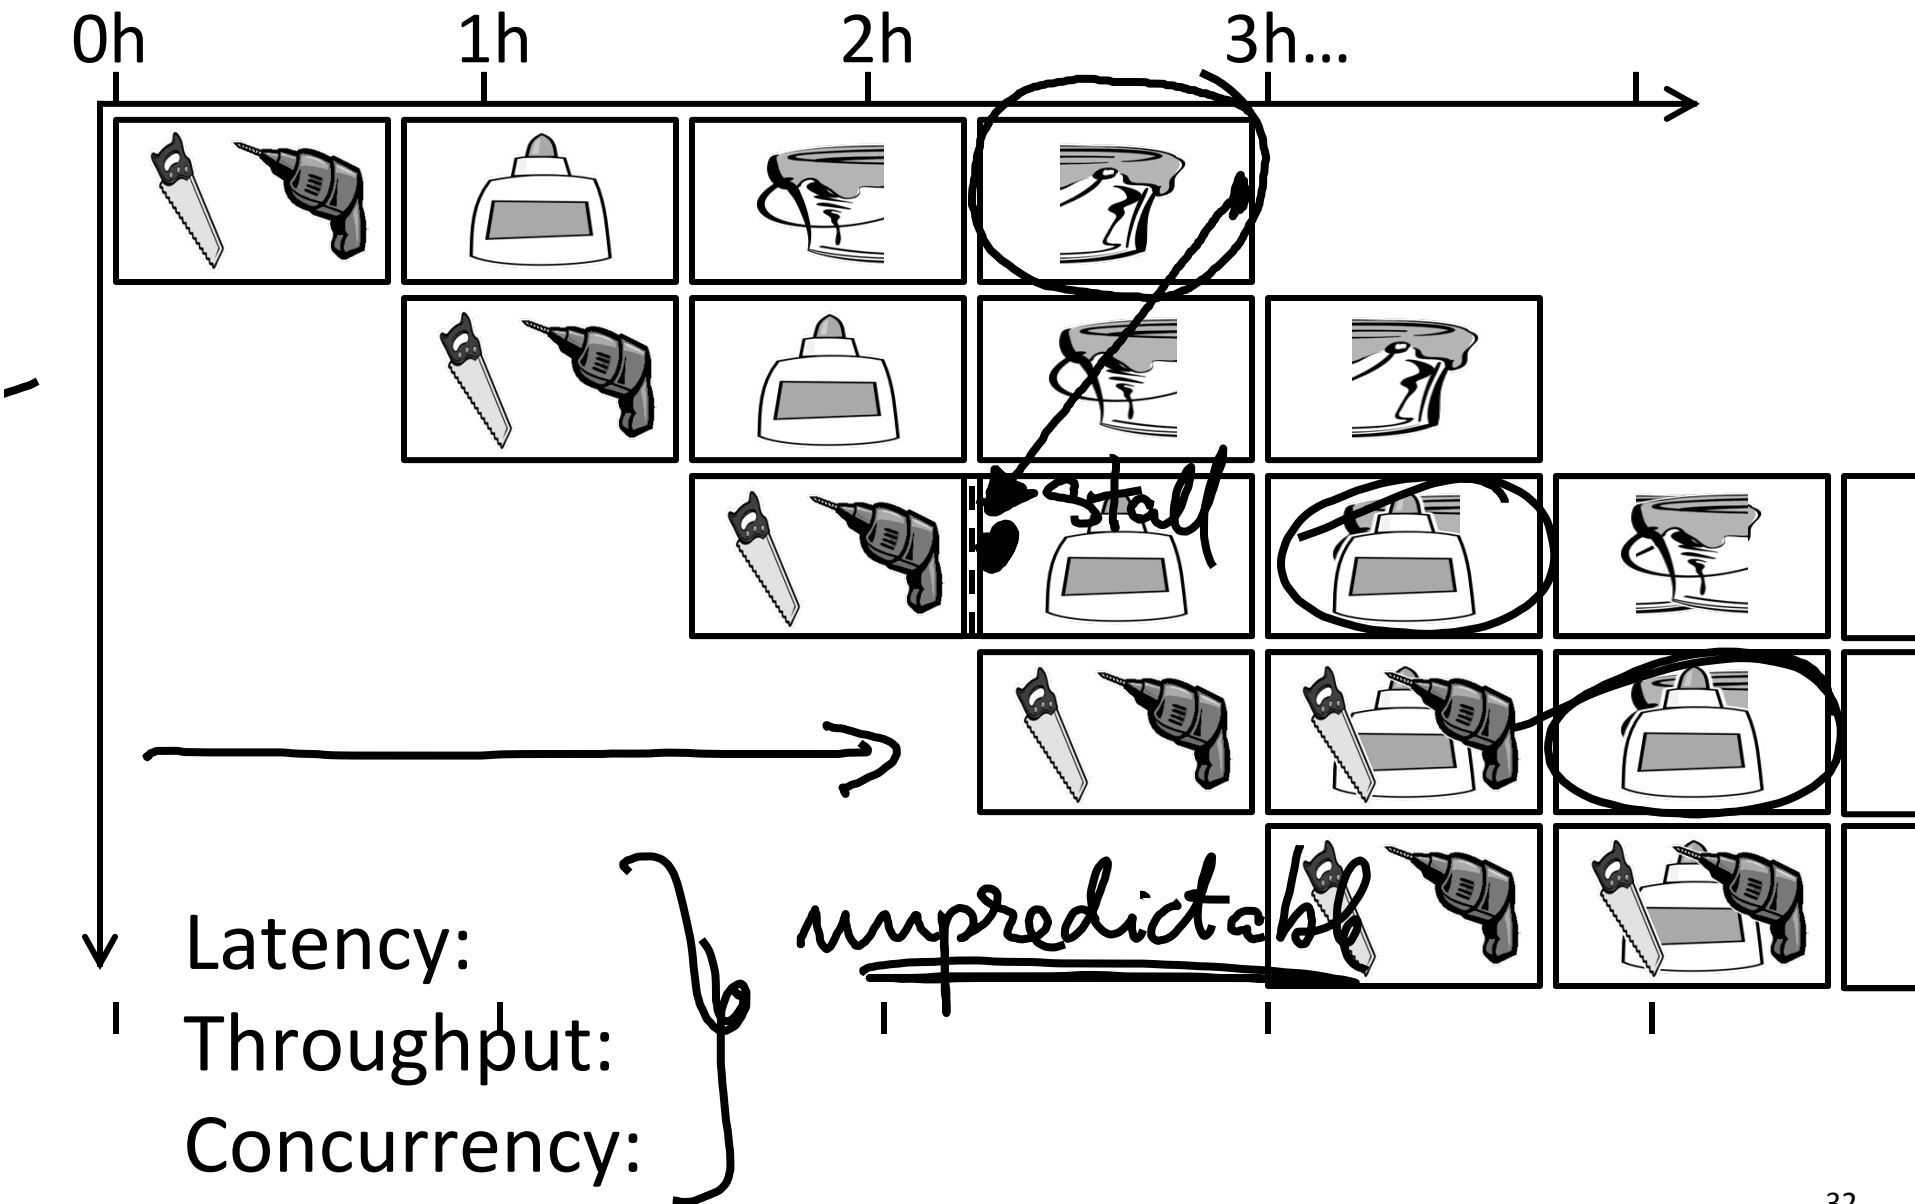
$$4$$





Latency: 3 hrs/task  
 Throughput: 1 task 45 min  
 Concurrency: 4

Q: What if glue step of task 3 depends on output of task 1?






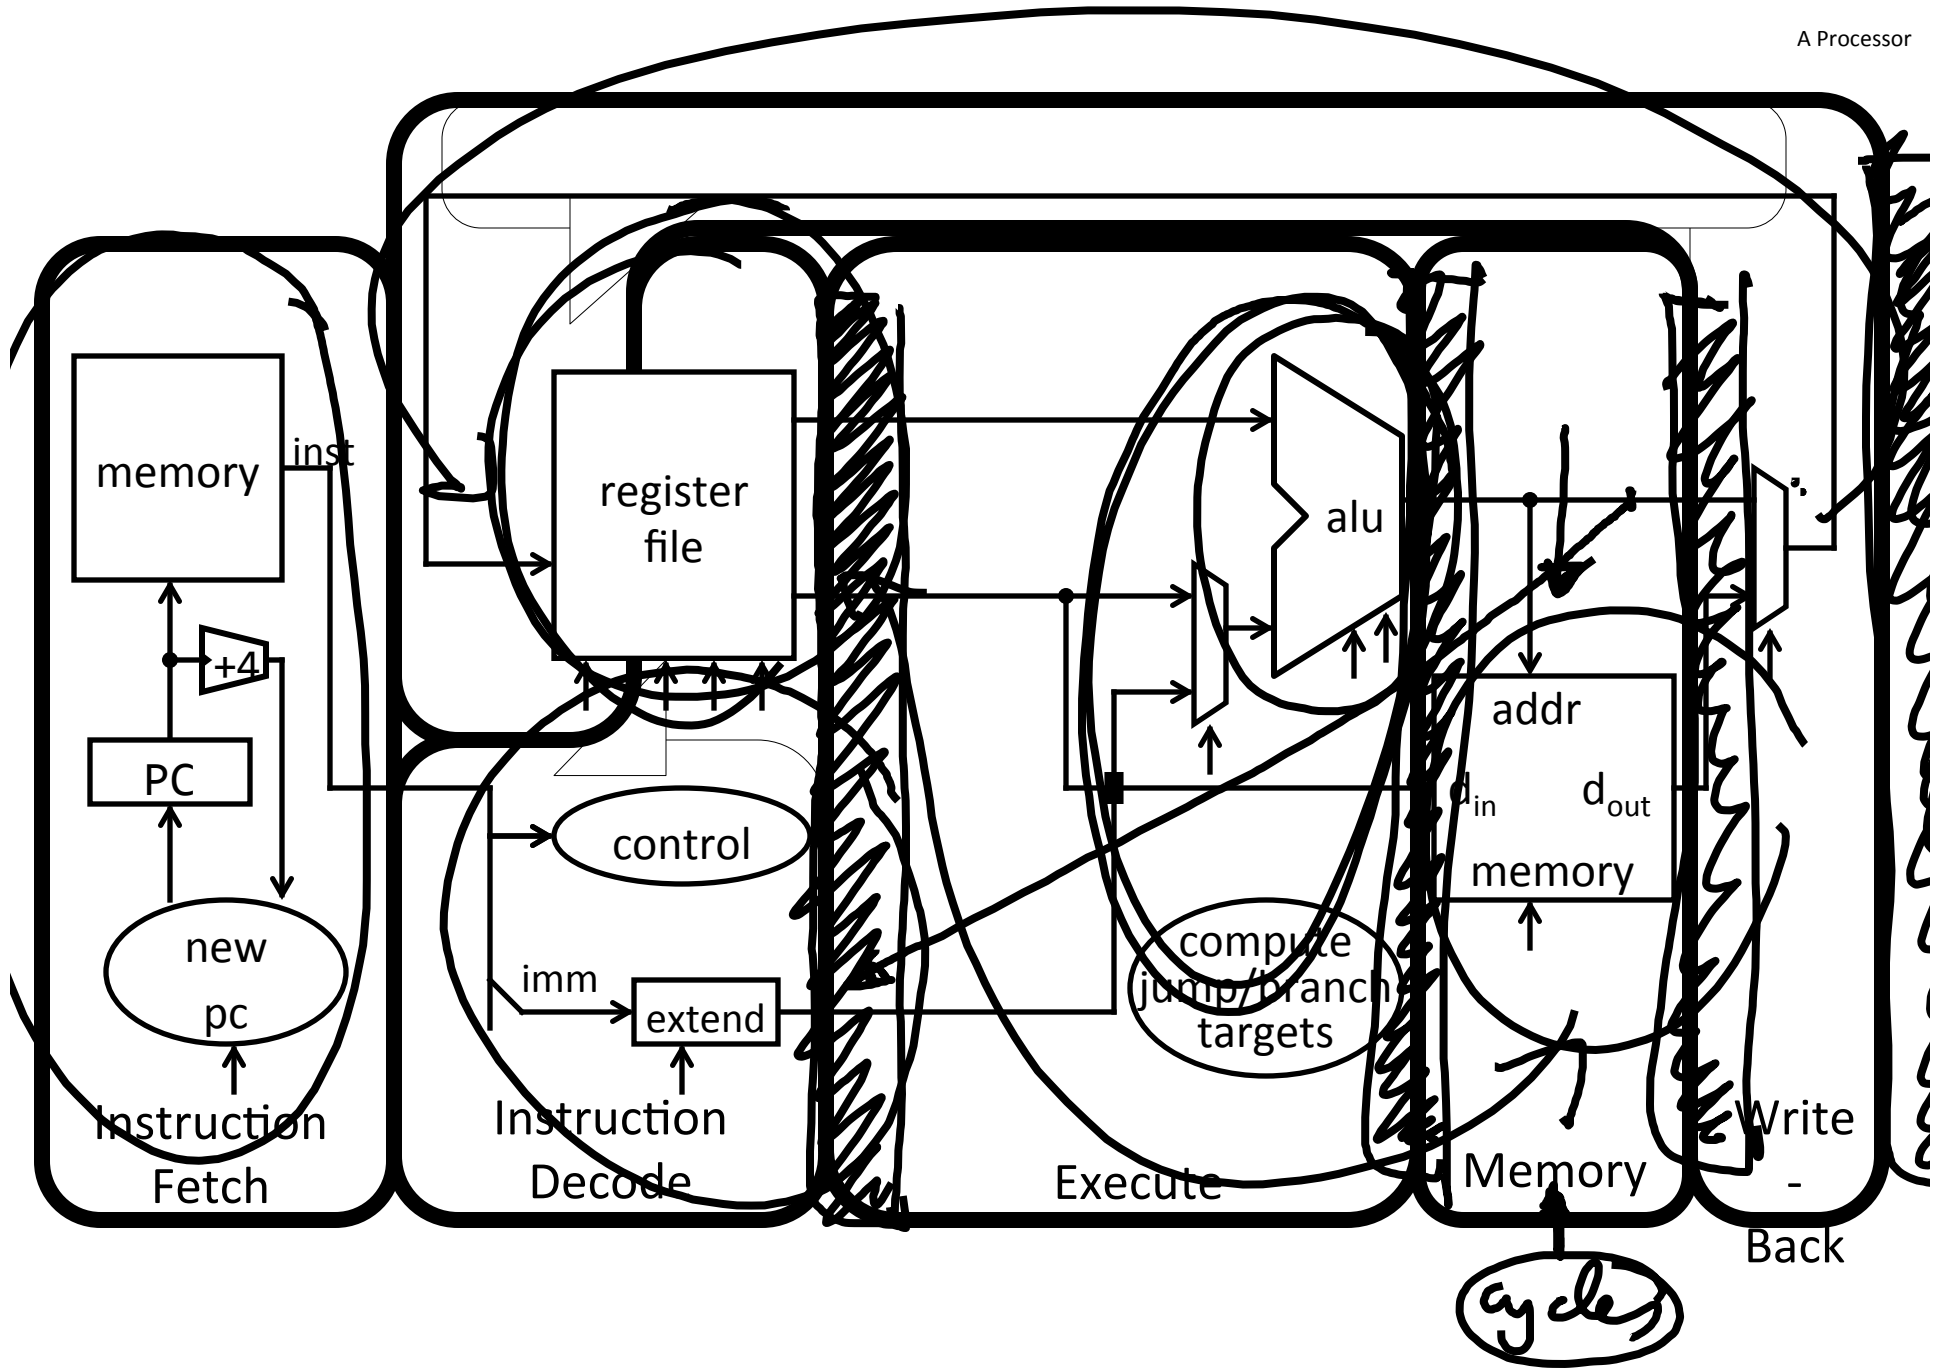
## Principle:

Throughput increased by parallel execution

## Pipelining:

- Identify *pipeline stages*
- Isolate stages from each other
- Resolve pipeline *hazards*





# Five stage "RISC" load-store architecture

## 1. Instruction fetch (IF)

- get instruction from memory, increment PC

## 2. Instruction Decode (ID)

- translate opcode into control signals and read registers

## 3. Execute (EX)

- perform ALU operation, compute jump/branch targets

## 4. Memory (MEM)

- access memory if needed

## 5. Writeback (WB)

- update register file

5 instr / cycle

(1 gate)

Break instructions across multiple clock cycles  
(five, in this case)

Design a separate stage for the execution  
performed during each clock cycle

Add pipeline registers (flip-flops) to isolate signals  
between different stages