

CPU Performance

Pipelined CPU

Hakim Weatherspoon
CS 3410, Spring 2012
Computer Science
Cornell University

See P&H Chapters 1.4 and 4.5

A Simple CPU: remaining branch instructions

Memory Layout

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)

	0x00000000
	0x00000001
	0x00000002
	0x00000003
	0x00000004
	0x00000005
	0x00000006
	0x00000007
	0x00000008
	0x00000009
	0x0000000a
	0x0000000b
	...
	0xffffffff

Control Flow: More Branches

Conditional Jumps (cont.)

000001001010000100000000000000010

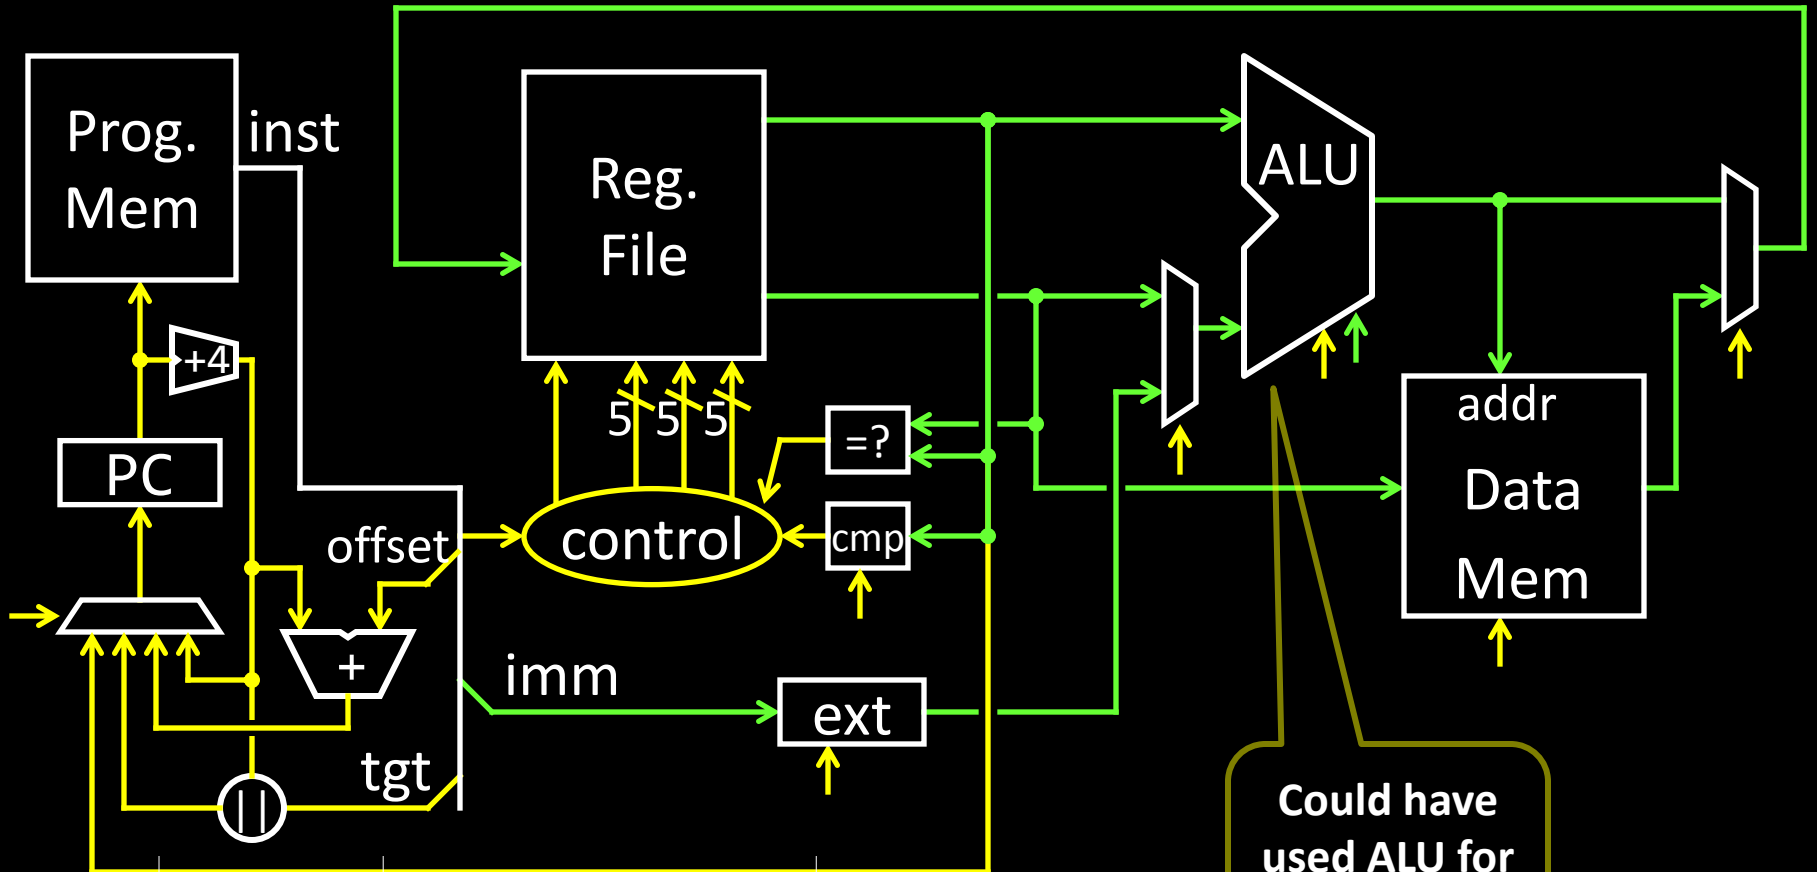
op
6 bits
rs
5 bits
subop
5 bits
offset
16 bits

almost I-Type

signed offsets

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 $PC = PC + 4 + (\text{offset} \ll 2)$
0x6	0x0	BLEZ rs, offset	if $R[rs] \leq 0$ then $PC = PC + 4 + (\text{offset} \ll 2)$
0x7	0x0	BGTZ rs, offset	if $R[rs] > 0$ then $PC = PC + 4 + (\text{offset} \ll 2)$

Absolute Jump



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 $PC = PC + 4 + (\text{offset} \ll 2)$
0x6	0x0	BLEZ rs, offset	if $R[rs] \leq 0$ then $PC = PC + 4 + (\text{offset} \ll 2)$
0x7	0x0	BGTZ rs, offset	if $R[rs] > 0$ then $PC = PC + 4 + (\text{offset} \ll 2)$

Control Flow: Jump and Link Function/procedure calls

00001100000001001000011000000010

op

6 bits

immediate

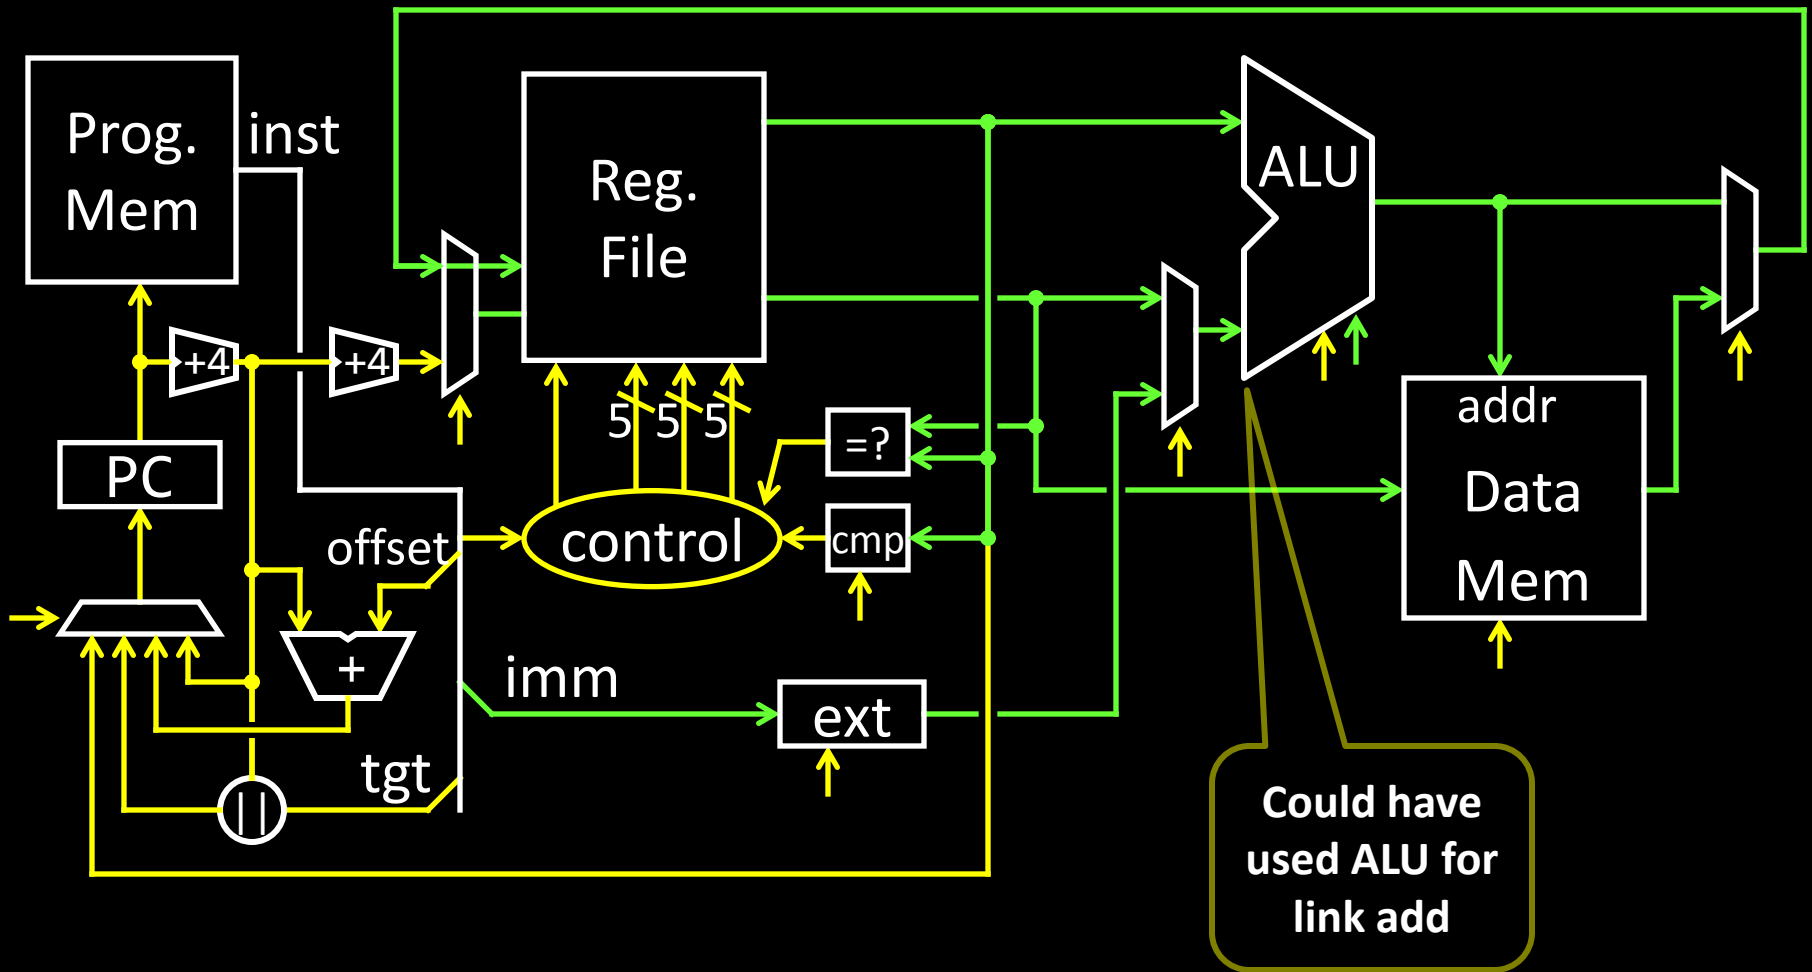
26 bits

J-Type

op	mnemonic	description
0x3	JAL target	r31 = PC+8 (+8 due to branch delay slot) PC = (PC+4) _{31..28} (target << 2)

op	mnemonic	description
0x2	J target	PC = (PC+4) _{31..28} (target << 2)

Absolute Jump



op	mnemonic	description
0x3	JAL target	$r31 = PC + 8$ (+8 due to branch delay slot) $PC = (PC + 4)_{31..28} (target \ll 2)$

Performance

See: P&H 1.4

Design Goals

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, ...

Performance

How to measure performance?

GHz (billions of cycles per second)

MIPS (millions of instructions per second)

MFLOPS (millions of floating point operations per second)

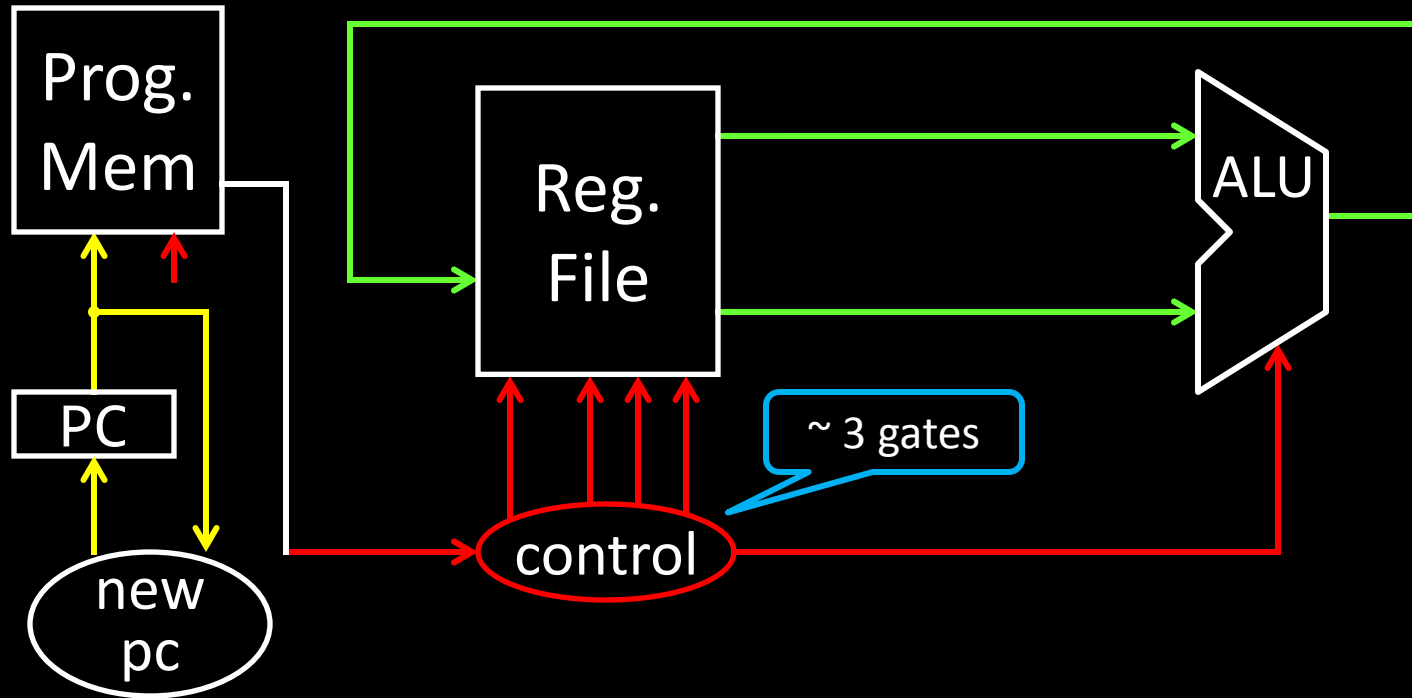
benchmarks (SPEC, TPC, ...)

Metrics

latency: how long to finish my program

throughput: how much work finished per unit time

How Fast?



Assumptions:

- alu: 32 bit ripple carry + some muxes
- next PC: 30 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

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

Adder Performance

32 Bit Adder Design	Space	Time
Ripple Carry	≈ 300 gates	≈ 64 gate delays
2-Way Carry-Skip	≈ 360 gates	≈ 35 gate delays
3-Way Carry-Skip	≈ 500 gates	≈ 22 gate delays
4-Way Carry-Skip	≈ 600 gates	≈ 18 gate delays
2-Way Look-Ahead	≈ 550 gates	≈ 16 gate delays
Split Look-Ahead	≈ 800 gates	≈ 10 gate delays
Full Look-Ahead	≈ 1200 gates	≈ 5 gate delays

Optimization: Summary

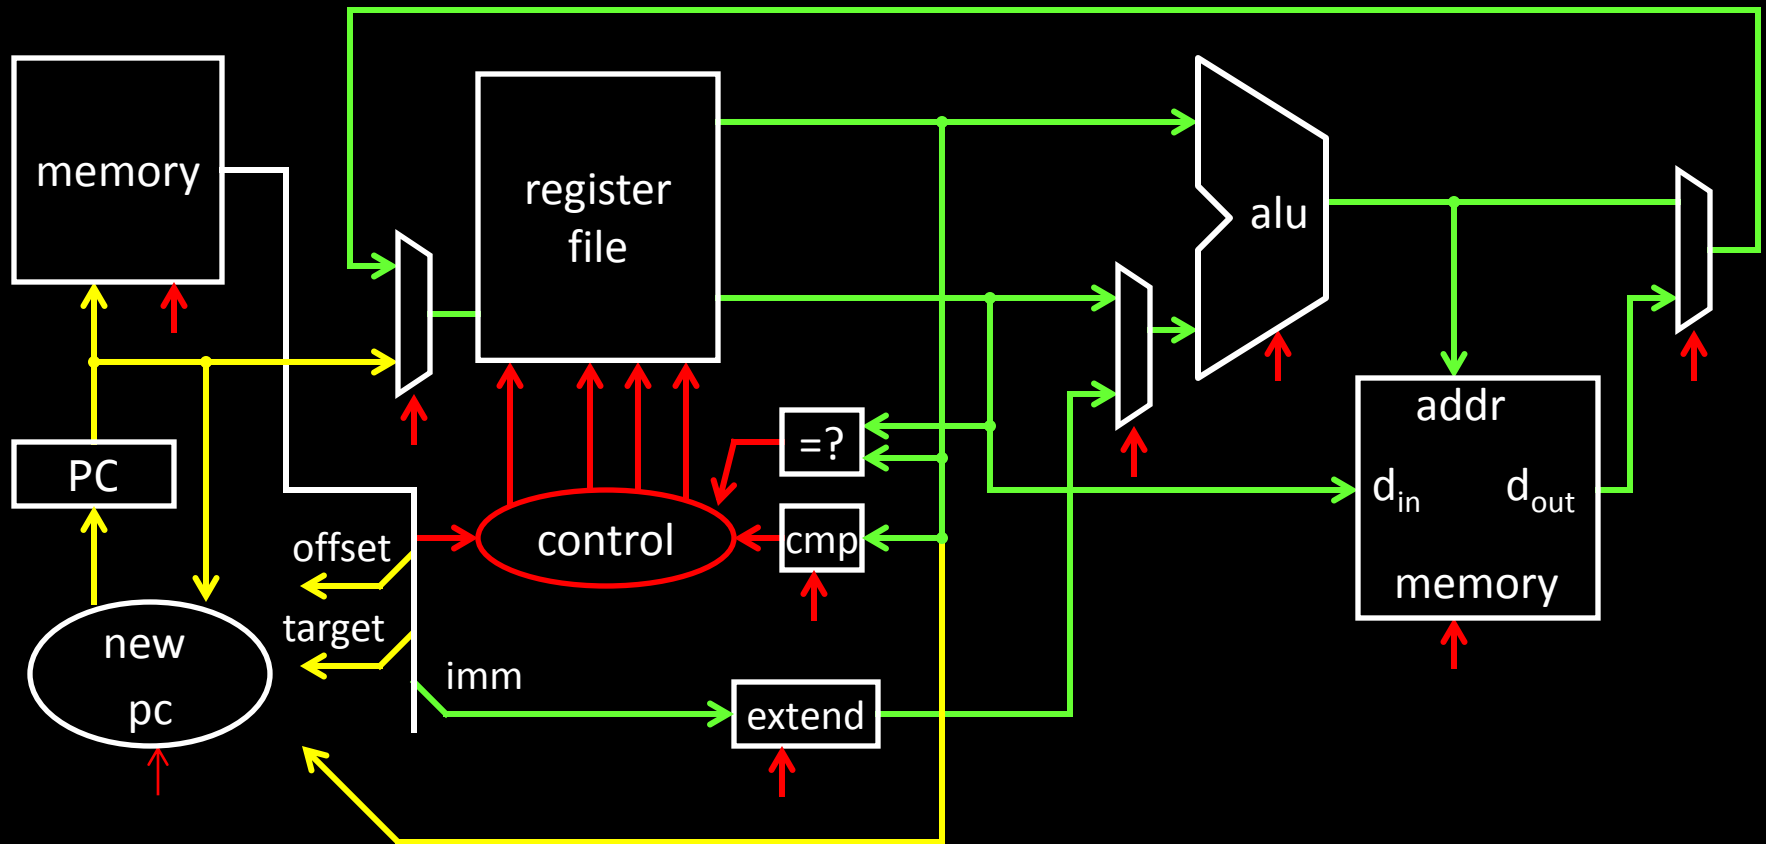
Critical Path

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

Strategy 1 (we just employed)

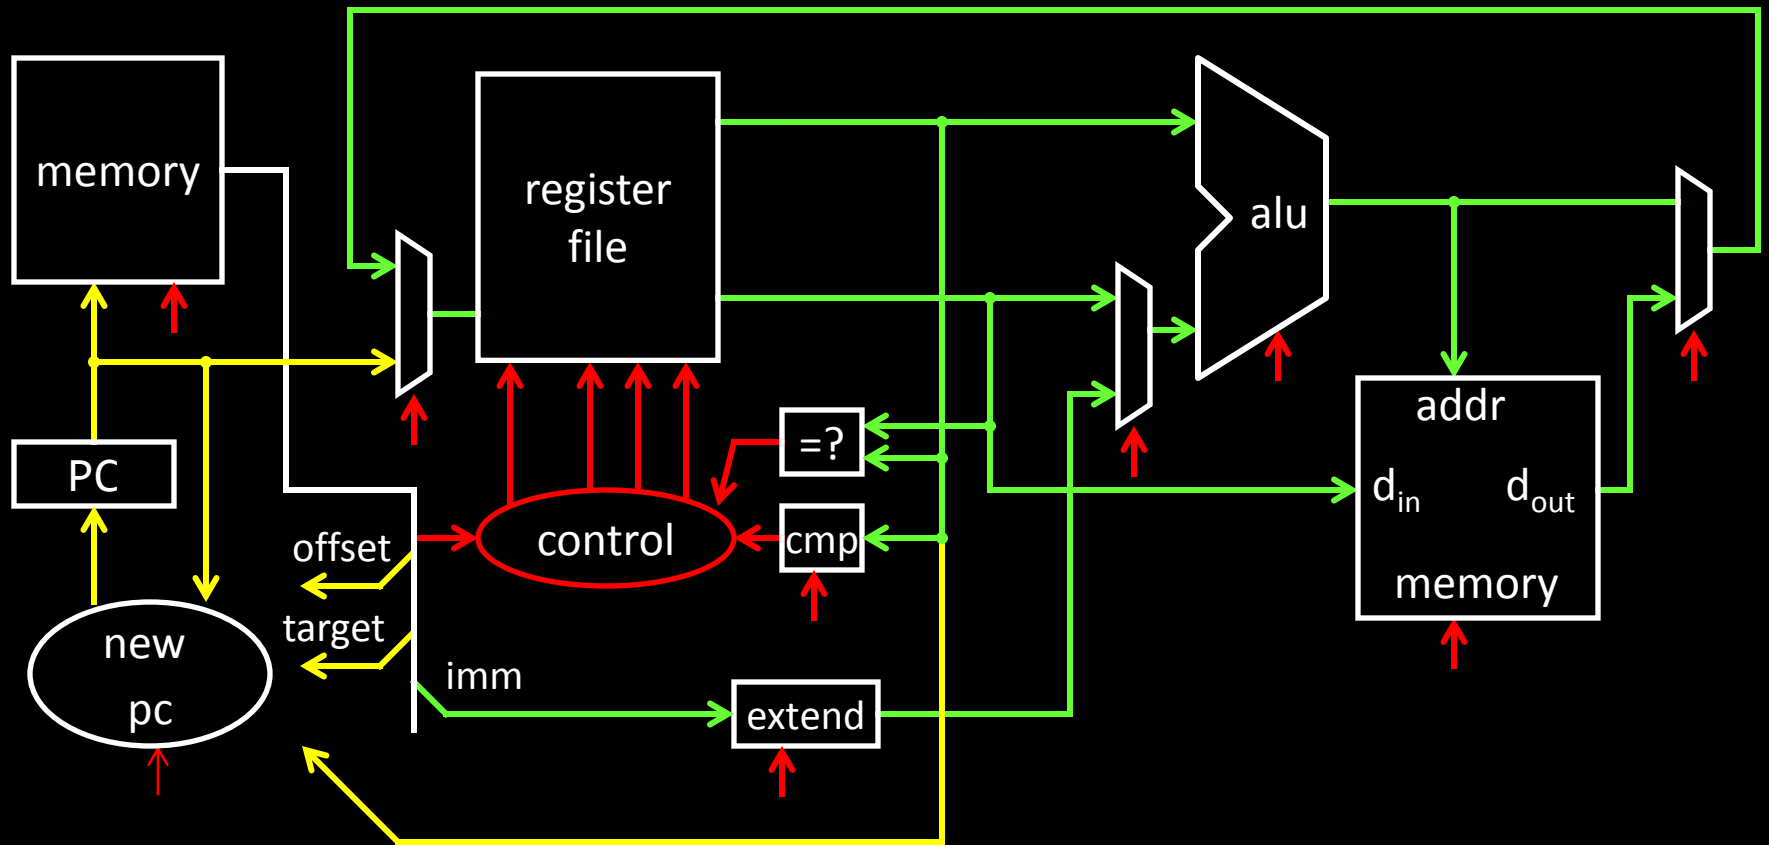
- Optimize for delay on the critical path
- Optimize for size / power / simplicity elsewhere
 - next PC

Processor Clock 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)	$\text{Mem}[\text{offset}+R[rs]] = R[rd]$
0x2b	SW rd, offset(rs)	$\text{Mem}[\text{offset}+R[rs]] = R[rd]$

Processor Clock Cycle



op	func	mnemonic	description
0x0	0x08	JR rs	PC = R[rs]

op	mnemonic	description
0x2	J target	PC = (PC+4) _{31..28} (target << 2)

Multi-Cycle Instructions

Strategy 2

- Multiple cycles to complete a single instruction

E.g: Assume:

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

Multi-Cycle CPU

30 MHz (33 ns cycle) with

- 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

CPI

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:

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

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

Example

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
- 15% branches, CPI = 1

Amdahl's Law

Amdahl's Law

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

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

The Kids

Alice



Bob



They don't always get along...

The Bicycle



The Materials



Saw



Drill




Glue




Paint

The Instructions

N pieces, each built following same sequence:



Design 1: Sequential Schedule




Alice owns the room

Bob can enter when Alice is finished

Repeat for remaining tasks

No possibility for conflicts

Sequential Performance



Latency:

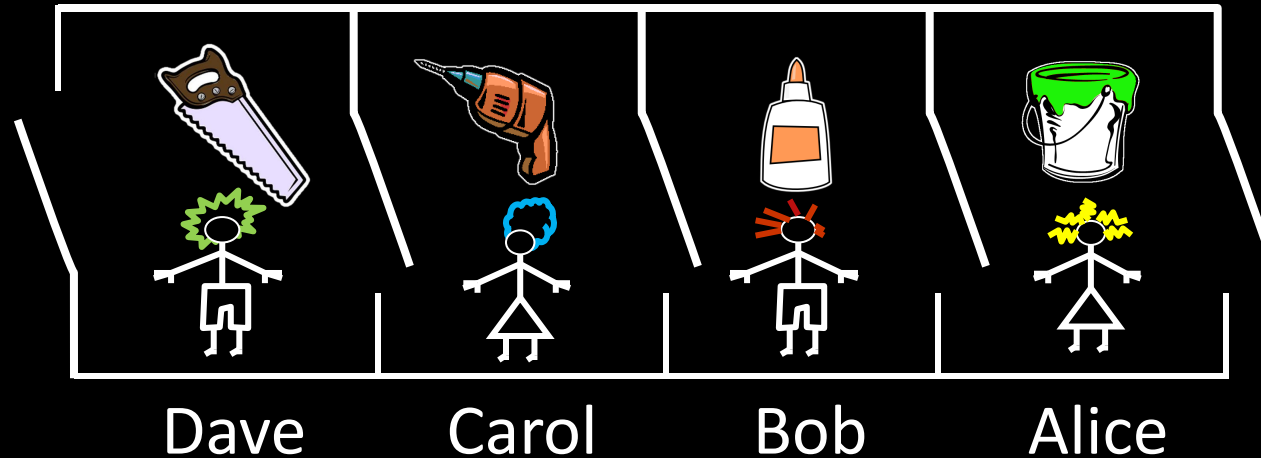
Throughput:

Concurrency:

Can we do better?

Design 2: Pipelined Design

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

time Pipelined Performance



Latency:

Throughput:

Concurrency:

Unequal Pipeline Stages



Latency:

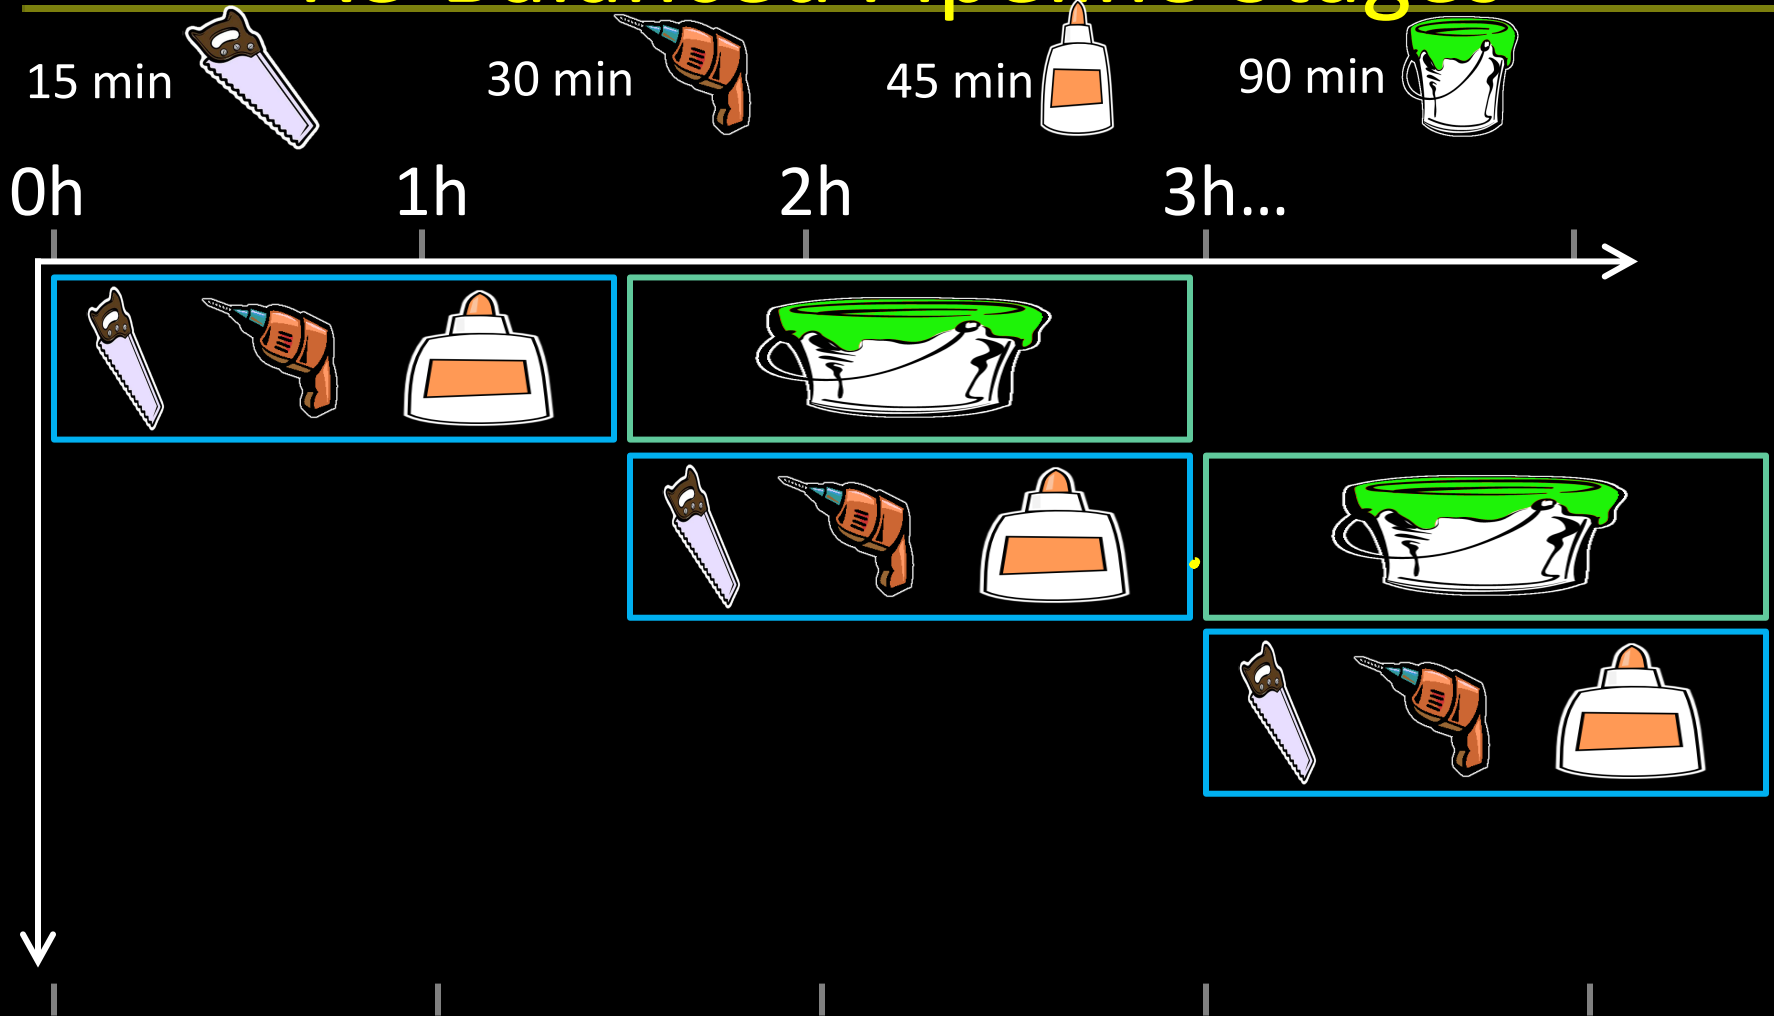
Throughput:

Concurrency:


Poorly-balanced Pipeline Stages



Re-Balanced Pipeline Stages



Splitting Pipeline Stages



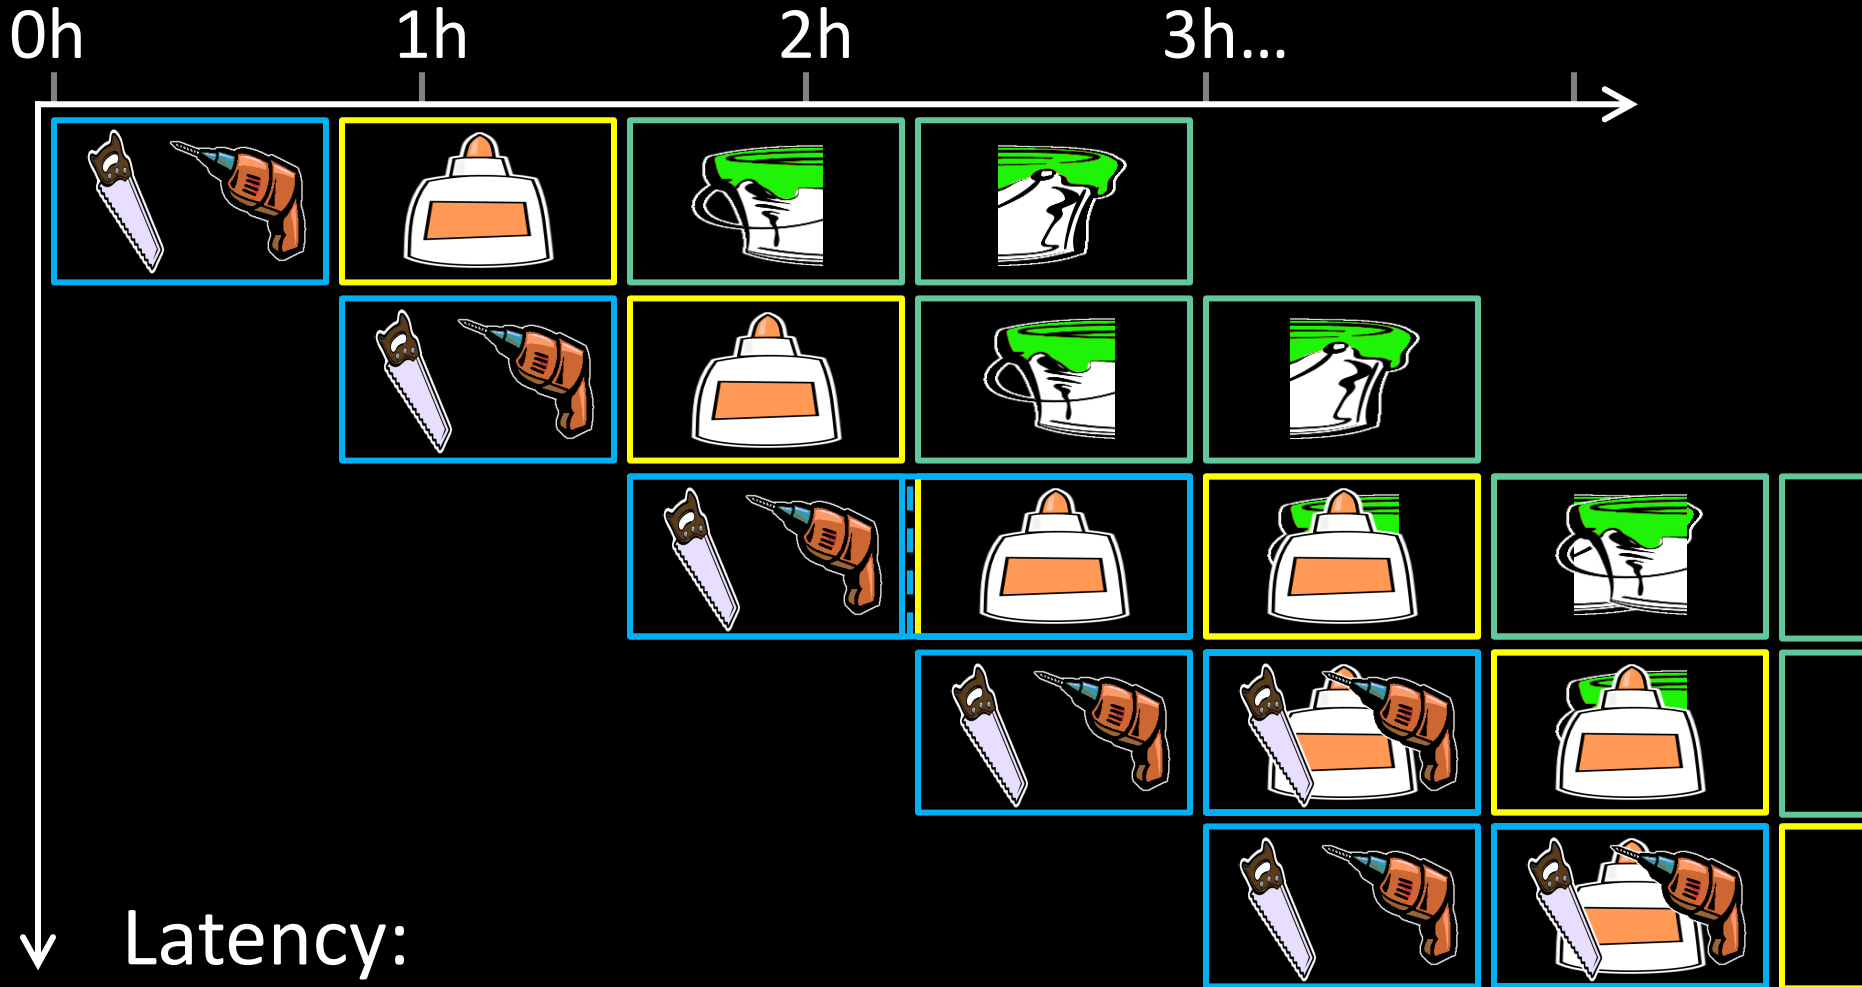
Latency:

Throughput:

Concurrency:

Pipeline Hazards

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



Latency:

Throughput:

Concurrency:

Lessons


Principle:

Throughput increased by parallel execution

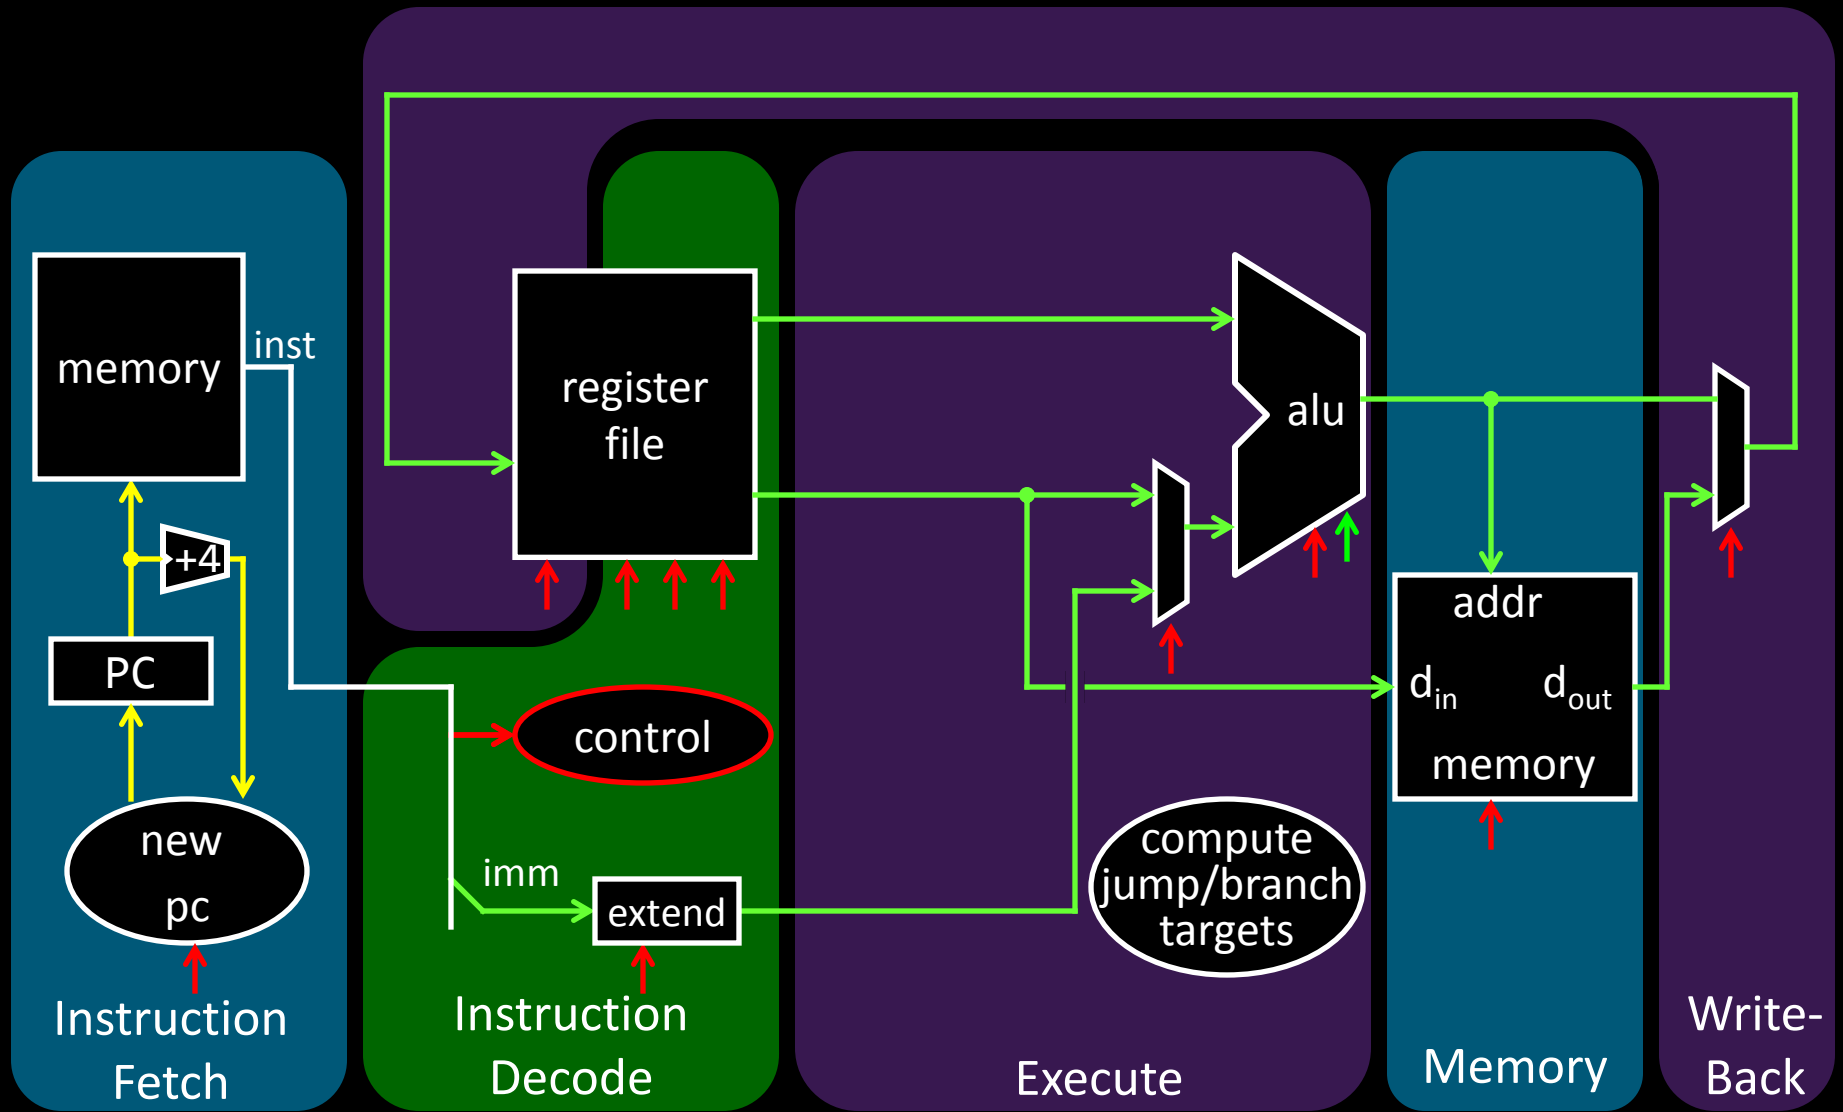
Pipelining:

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

A Processor



A Processor



Basic Pipeline

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

Pipelined Implementation

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