

Midterms on the table

Lecture 18: Pipelining



Last names: A-D E-L M-R S-Z
(for the most part)

- Today's topics:

- Power and energy
- 5-stage pipeline
- Hazards
- Data dependence handling with bypassing
- Data dependence examples

Email me if your midterm
score on Canvas is zero!!

HW7 posted
later today

Midterm Notes

240 43 45

- Grade categories:

- Top 18%: 93 (A)
- Top 36%: 87 (A-)
- Top 54%: 81 (B+)
- Top 72%: 70 (B)
- Top 88%: 60 (B-)

A
—— 93
A-
—— 87
B+
—— 81
B
—— 70
B-
—— 60
○

- Common mistakes:

- Support for subtraction: not including the carry
- IEEE 754 floating point formats and addition
- Addition of signed numbers
- Logic gates, Sum-of-products ←
- MARS variables, syscalls ←
- Power and energy !!! ←

\$a0 ←
\$v0 ←
syscall

Power and Energy

$$\text{Total Power} = \text{Dyn Power} + \text{Lkg power}$$

$$\text{Lkg power} \propto \text{Voltage} \rightarrow \text{Lkg stays constant at 20W (in midten question)}$$

$$\text{Dyn power} \propto \text{Act} \times \text{C} \times \text{f} \times \text{V}^2$$

$$100\text{W} = 80\text{W dyn} + 20\text{W lkg} \Rightarrow 2\text{GHz}$$

$$\underline{140\text{W}} = 80\text{W} \times \frac{3}{2} + 20\text{W} \Rightarrow 3\text{GHz Turbo}$$

$$\underline{60\text{W}} = 80 \times \frac{1}{2} + 20\text{W} \Rightarrow 1\text{GHz}$$

Power \Rightarrow Watts

Program takes 10 secs \Rightarrow CPU-bound

$$\begin{aligned} \text{At } 3\text{GHz, new exec time} &= 10 \times \frac{2}{3} = 6.66\text{secs} \\ \text{At } 2\text{GHz, new exec time} &= 10 \times \frac{2}{1} = 20\text{secs} \end{aligned}$$

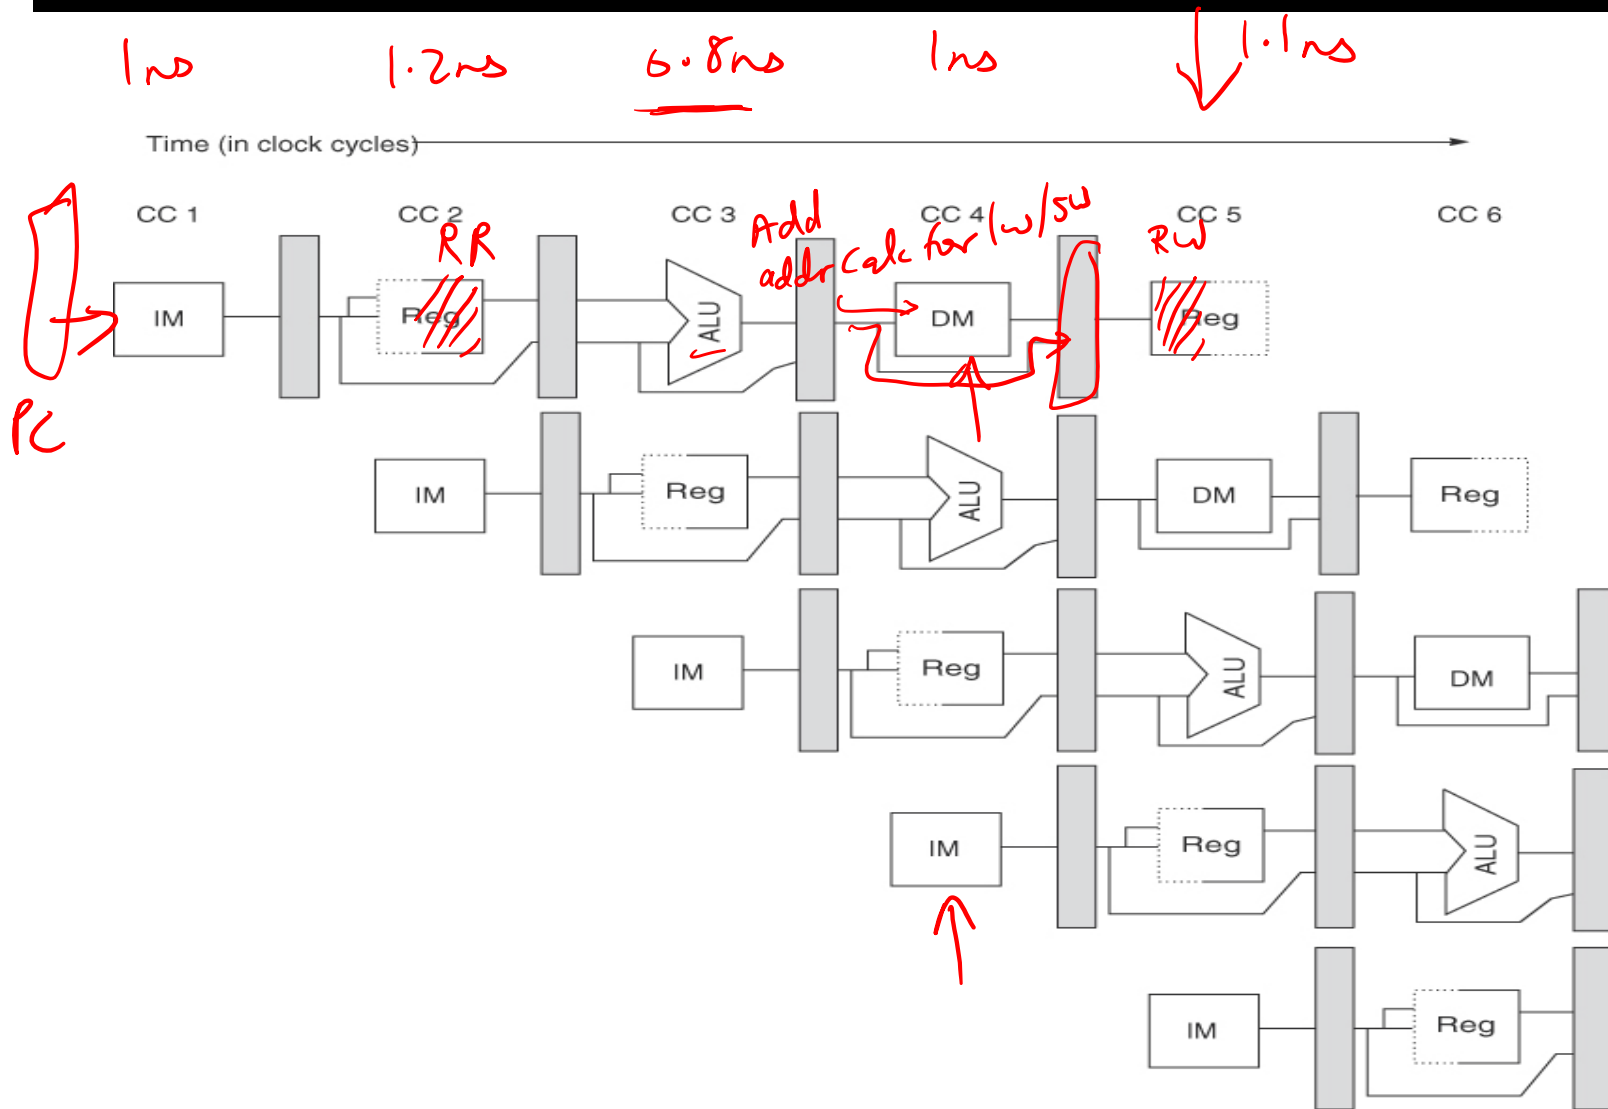
$$E = P \times t$$
$$E = 140 \times 6.66 = 920$$
$$E = 60 \times 20 = 1200 \text{ J}$$

3 Joules

Power and Energy

A 5-Stage Pipeline

Cycle time = 1.2ns



Performance Improvements?

- Does it take longer to finish each individual job? *Yes*
- Does it take shorter to finish a series of jobs? *Yes, before of parallelism*
- What assumptions were made while answering these questions?

- No dependences between instructions
- Easy to partition circuits into uniform pipeline stages
- No latch overhead


latches are an additional overhead

- Is a 10-stage pipeline better than a 5-stage pipeline?

10x impr

5x impr

ideal condts



Quantitative Effects

Going from unpipelined to pipelined


- As a result of pipelining:

- Sound & bad {
- Time in ns per instruction goes up *becoz of latch overhead*
 - Each instruction takes more cycles to execute *defn of cycle has changed*
 - But... average CPI remains roughly the same
 - ✓ ➤ Clock speed goes up
 - ✓ ➤ Total execution time goes down, resulting in lower average time per instruction
 - Under ideal conditions, speedup
= ratio of *elapsed times between successive instruction completions*
= number of pipeline stages = increase in clock speed

Hazards

- Structural hazards: different instructions in different stages (or the same stage) conflicting for the same resource
- Data hazards: an instruction cannot continue because it needs a value that has not yet been generated by an earlier instruction
- Control hazard: fetch cannot continue because it does not know the outcome of an earlier branch – special case of a data hazard – separate category because they are treated in different ways

Conflicts/Problems



- ^{IM} I-cache and ^{DM} D-cache are accessed in the same cycle – it helps to implement them separately
- Registers are read and written in the same cycle – easy to deal with if register read/write time equals cycle time/2 ✓
- Instructions can't skip the DM stage, else conflict for RW
- Consuming instruction may have to wait for producer
- Branch target changes only at the end of the second stage
-- what do you do in the meantime?

Structural Hazards

- Example: a unified instruction and data cache → stage 4 (MEM) and stage 1 (IF) can never coincide
- The later instruction and all its successors are delayed until a cycle is found when the resource is free → these are pipeline bubbles
- Structural hazards are easy to eliminate – increase the number of resources (for example, implement a separate instruction and data cache, add more register ports)

Data Hazards

prod \$2 ← ~~to~~
~~to~~ ← \$2 consumer
~~to~~

- An instruction *produces* a value in a given pipeline stage
- A subsequent instruction *consumes* that value in a pipeline stage
- The consumer may have to be delayed so that the time of consumption is later than the time of production

Decode $\xrightarrow{D/R}$ RegRd
Example 1 – No Bypassing

I1 : IF — D/R — ALU — DM — RW
 $\begin{matrix} I1 \\ I3 \\ I2 \end{matrix}$ \longrightarrow time

- Show the instruction occupying each stage in each cycle (no bypassing)
 if I1 is $R1+R2 \rightarrow R3$ and I2 is $R3+R4 \rightarrow R5$ and I3 is $R7+R8 \rightarrow R9$

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

IF $I1$	IF $I2$	IF $I3$	IF $I3$	IF $I3$	IF	IF	IF
D/R	D/R $I1$	D/R $I2$	D/R $I2$	D/R $I2$	D/R $I3$	D/R	D/R
ALU	ALU	ALU $I1$	ALU ○	ALU ○	ALU $I2$	ALU $I3$	ALU
DM	DM	DM	DM $I1$	DM ○	DM ○	DM $I2$	DM $I3$
RW	RW	RW	RW	RW $I1$	RW ○	RW ○	RW $I2$

$$CPI = \frac{5}{3} = 1.66$$

$$IPC = \frac{3}{5} = 0.6$$

12 RW
 $I3$

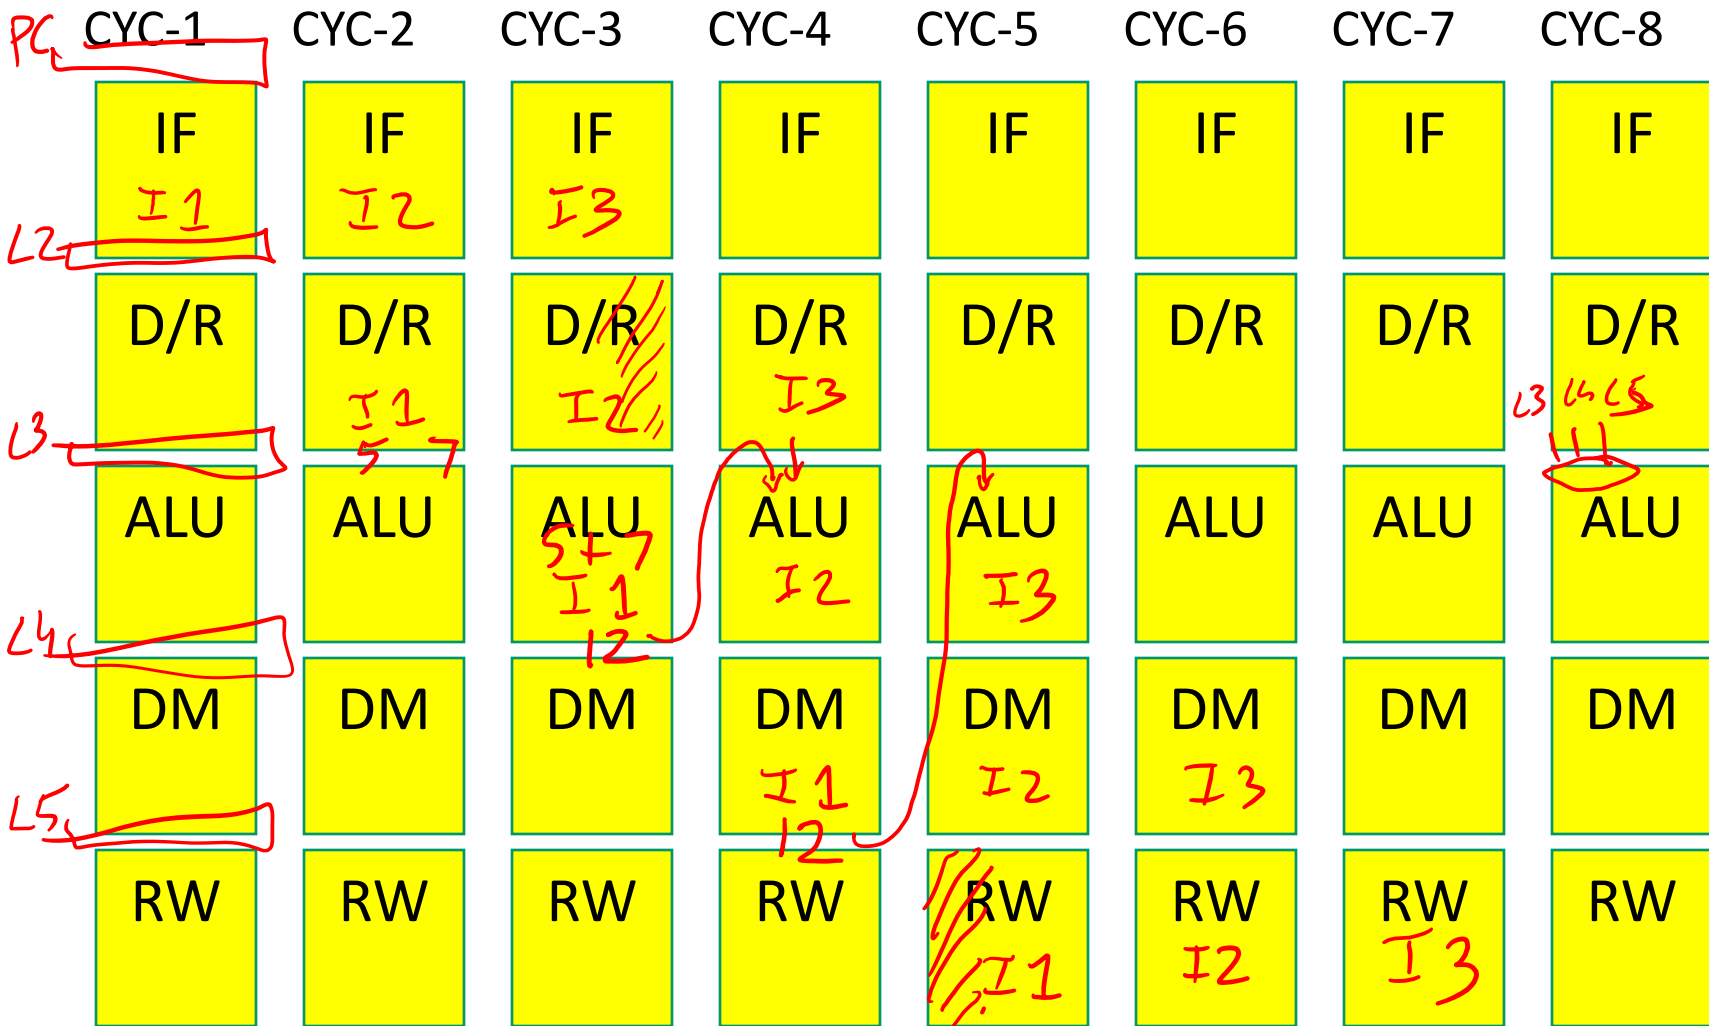
Example 1 – No Bypassing

- Show the instruction occupying each stage in each cycle (no bypassing)
if I1 is $R1+R2 \rightarrow R3$ and I2 is $R3+R4 \rightarrow R5$ and I3 is $R7+R8 \rightarrow R9$

CYC-1	CYC-2	CYC-3	CYC-4	CYC-5	CYC-6	CYC-7	CYC-8
IF I1	IF I2	IF I3	IF I3	IF I3	IF I4	IF I5	IF
D/R	D/R I1	D/R I2	D/R I2	D/R I2	D/R I3	D/R I4	D/R
ALU	ALU	ALU I1	ALU	ALU	ALU I2	ALU I3	ALU
DM	DM	DM	DM I1	DM	DM	DM I2	DM I3
RW	RW	RW	RW	RW I1	RW	RW	RW I2

Example 2 – Bypassing

- Show the instruction occupying each stage in each cycle (with bypassing)
if I1 is $R1+R2 \rightarrow R3$ and I2 is $R3+R4 \rightarrow R5$ and I3 is $R3+R8 \rightarrow R9$.
Identify the input latch for each input operand. 5 + 7 12 12 + 14




Example 2 – Bypassing

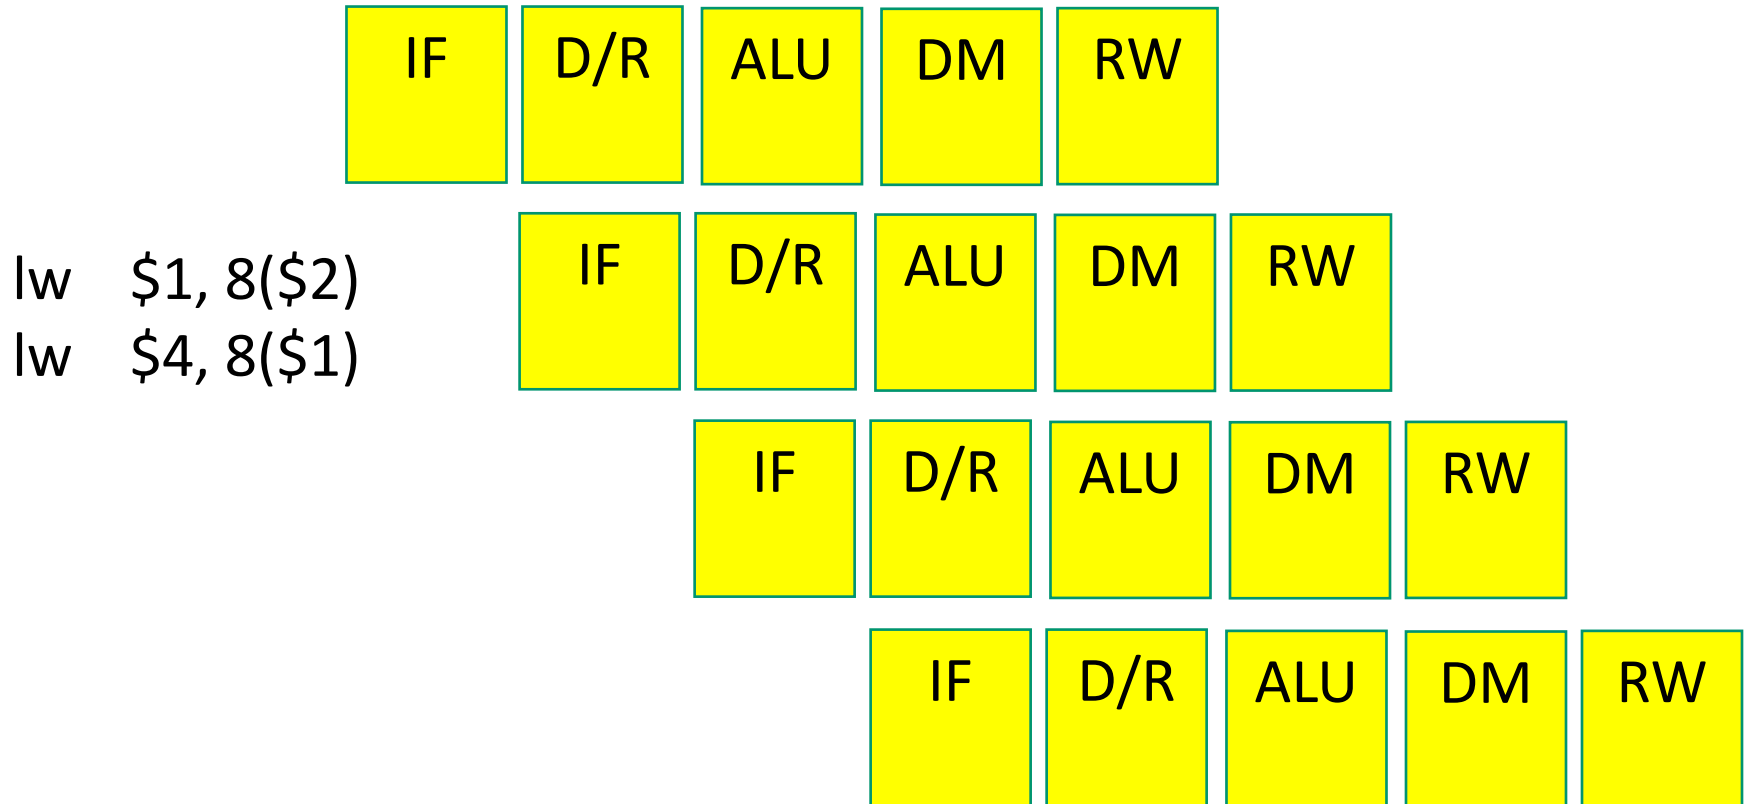
- Show the instruction occupying each stage in each cycle (with bypassing)
if I1 is $R1+R2 \rightarrow R3$ and I2 is $R3+R4 \rightarrow R5$ and I3 is $R3+R8 \rightarrow R9$.
Identify the input latch for each input operand.

CYC-1	CYC-2	CYC-3	CYC-4	CYC-5	CYC-6	CYC-7	CYC-8
IF I1	IF I2	IF I3	IF I4	IF I5	IF	IF	IF
D/R	D/R I1	D/R I2 I3	D/R I3 I4	D/R I4 I5	D/R	D/R	D/R
ALU	ALU	ALU I1	ALU I2	ALU I3	ALU	ALU	ALU
DM	DM	DM	DM I1	DM I2	DM I3	DM	DM
RW	RW	RW	RW	RW I1	RW I2	RW I3	RW

Problem 1



Problem 2



Problem 3

