


Embedded Systems



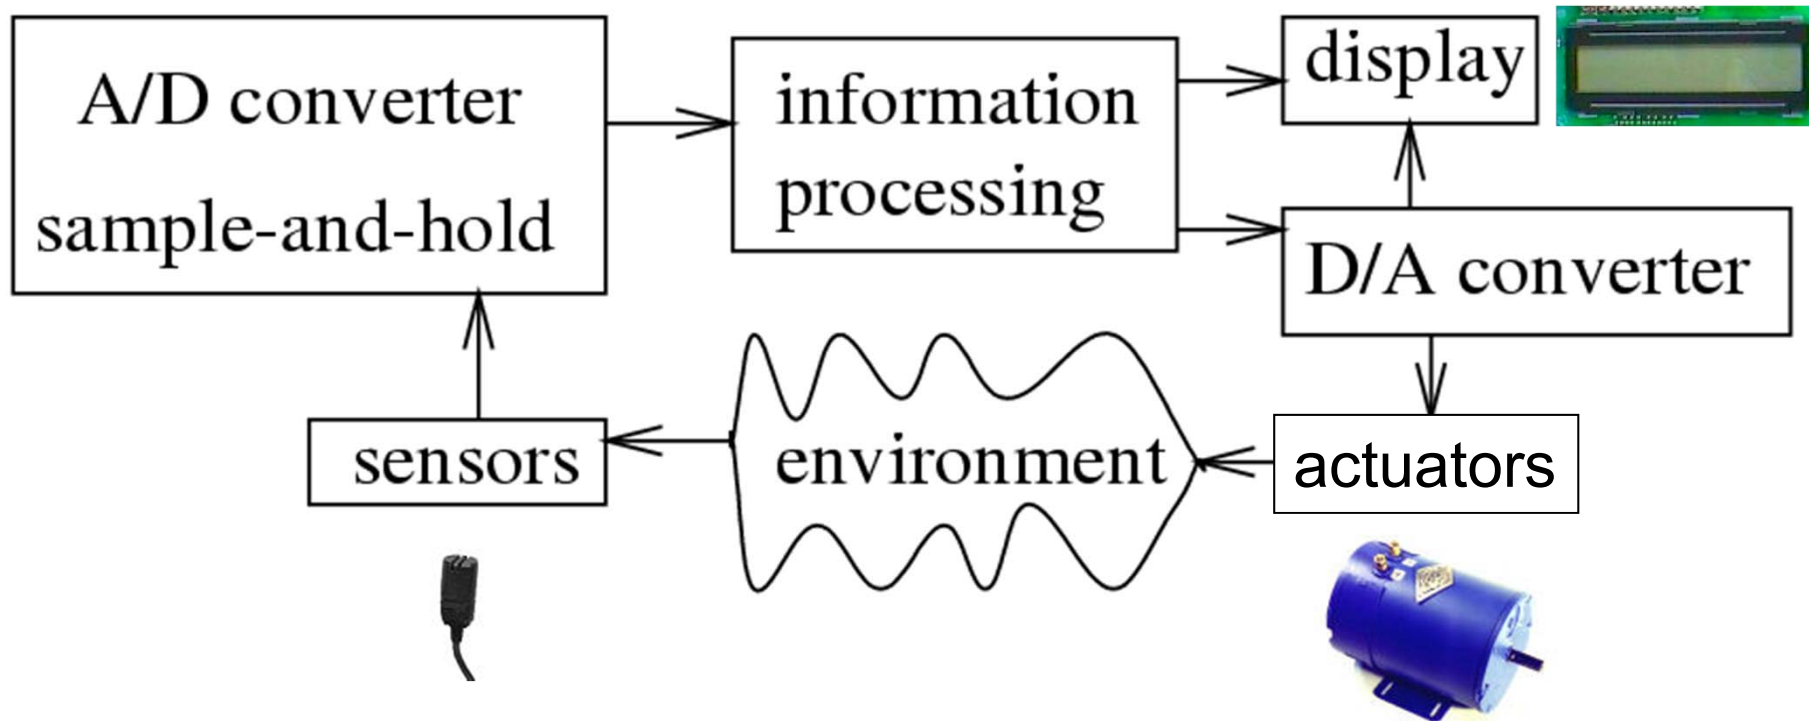


Dr. Eric Armengaud from the Virtual Vehicle Competence Center is going to give a talk on **model-based development and test of distributed automotive embedded systems** on **Tuesday, Jan. 11th**.

- Automotive embedded Systems
- SW Engineering
- networks (focus FlexRay)

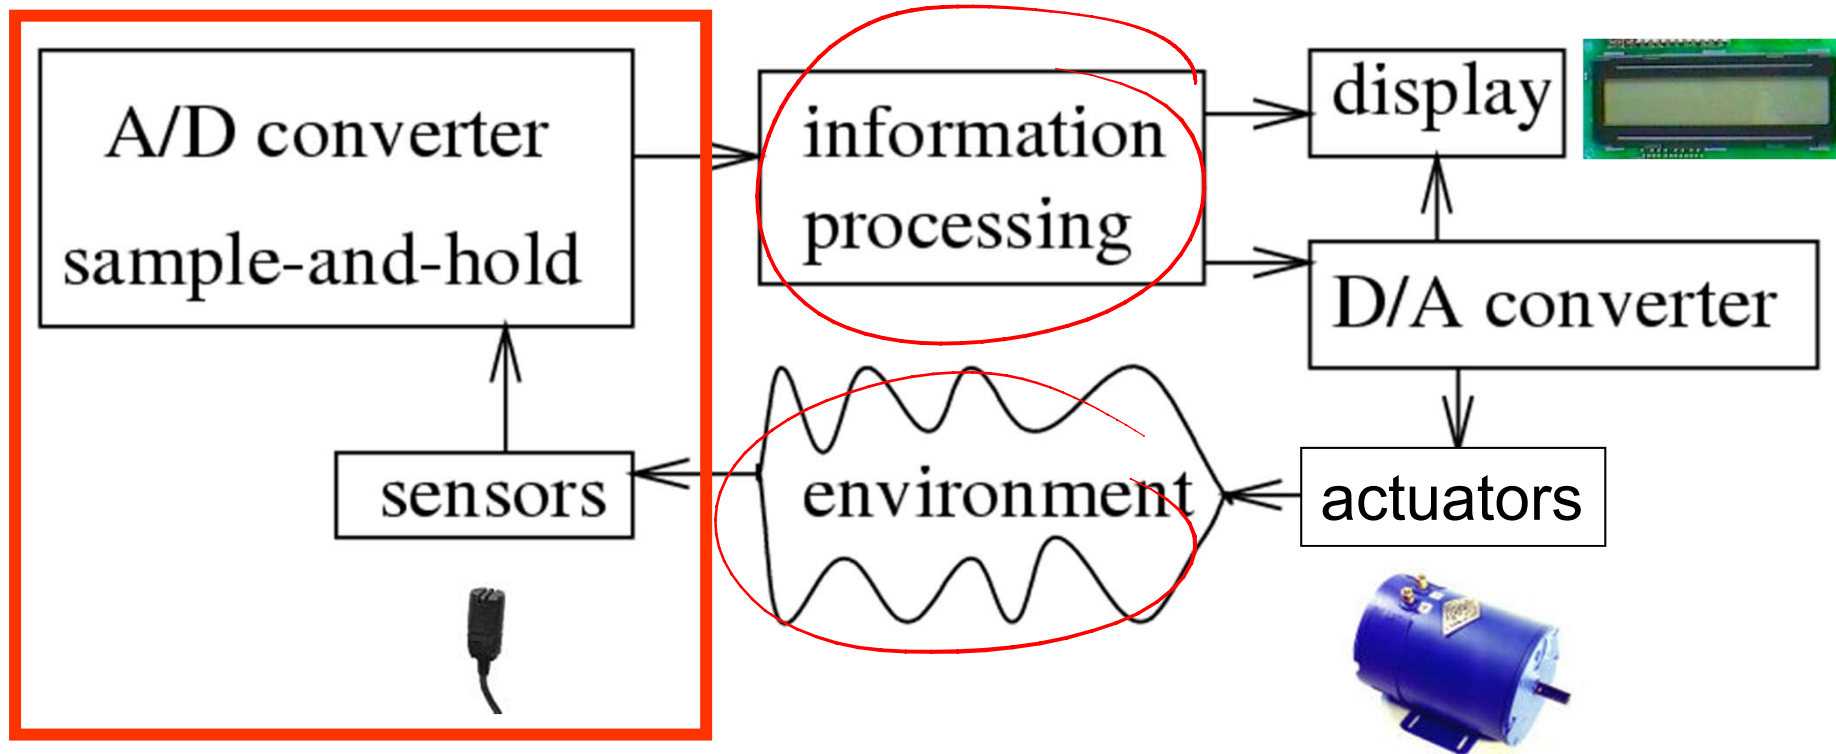
Embedded System Hardware

- Embedded system hardware is frequently used in a loop (*„hardware in a loop“*):



Embedded System Hardware

- Embedded system hardware is frequently used in a loop („*hardware in a loop*“):



TI Embedded Processing Portfolio

TI Embedded Processors

Microcontrollers (MCUs) ARM®-Based Processors Digital Signal Processors (DSPs)

16-bit ultra-low power MCUs 32-bit real-time MCUs 32-bit ARM Cortex™-M3 MCUs ARM Cortex-A8 MPUs DSP DSP+ARM Multi-core DSP Ultra Low power DSP


MSP430™

Up to 25 MHz

Flash
1 KB to 256 KB
Analog I/O, ADC
LCD, USB, RF

Measurement,
Sensing, General
Purpose

\$0.25 to \$9.00



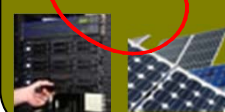
**C2000™
Delfino™
Piccolo™**

40MHz to 300 MHz

Flash, RAM
16 KB to 512 KB

PWM, ADC,
CAN, SPI, I²C
Motor Control,
Digital Power,
Lighting, Ren. Enrgy

\$1.50 to \$20.00



Stellaris®
ARM® Cortex™-M3


Up to 100 MHz

Flash
8 KB to 256 KB

USB, ENET
MAC+PHY CAN,
ADC, PWM, SPI

Connectivity, Security
Motion Control, HMI,
Industrial Automation

\$1.00 to \$8.00




Sitara™
ARM® Cortex™-A8
& ARM9

300MHz to >1GHz

Cache,
RAM, ROM
USB, CAN,
PCIe, EMAC

Industrial computing,
POS & portable
data terminals

\$5.00 to \$20.00



**C6000™
DaVinci™
video processors
OMAP™**


300MHz to >1Ghz
+Accelerator

Cache
RAM, ROM

USB, ENET,
PCIe, SATA, SPI

Floating/Fixed Point
Video, Audio, Voice,
Security, Confer.

\$5.00 to \$200.00



C6000™


24.000
MMACS

Cache
RAM, ROM

SRIO, EMAC
DMA, PCIe

Telecom T&M,
media gateways,
base stations

\$40 to \$200.00



C5000™


Up to 300 MHz
+Accelerator

Up to 320KB RAM
Up to 128KB ROM
USB, ADC
McBSP, SPI, I²C

Audio, Voice

Medical, Biometrics

\$3.00 to \$10.00







Software & Dev. Tools




Piccolo™ controlSTICK




Broad C2000 Application Base



**Renewable
Energy
Generation**

**Telecom
Digital Power**

**AC Drives, Industrial
& Consumer Motor
Control**



**Automotive
Radar, Electric
Power Steering
& Digital Power**



**Power Line
Communications**



**Consumer,
Medical &
Non-traditional**

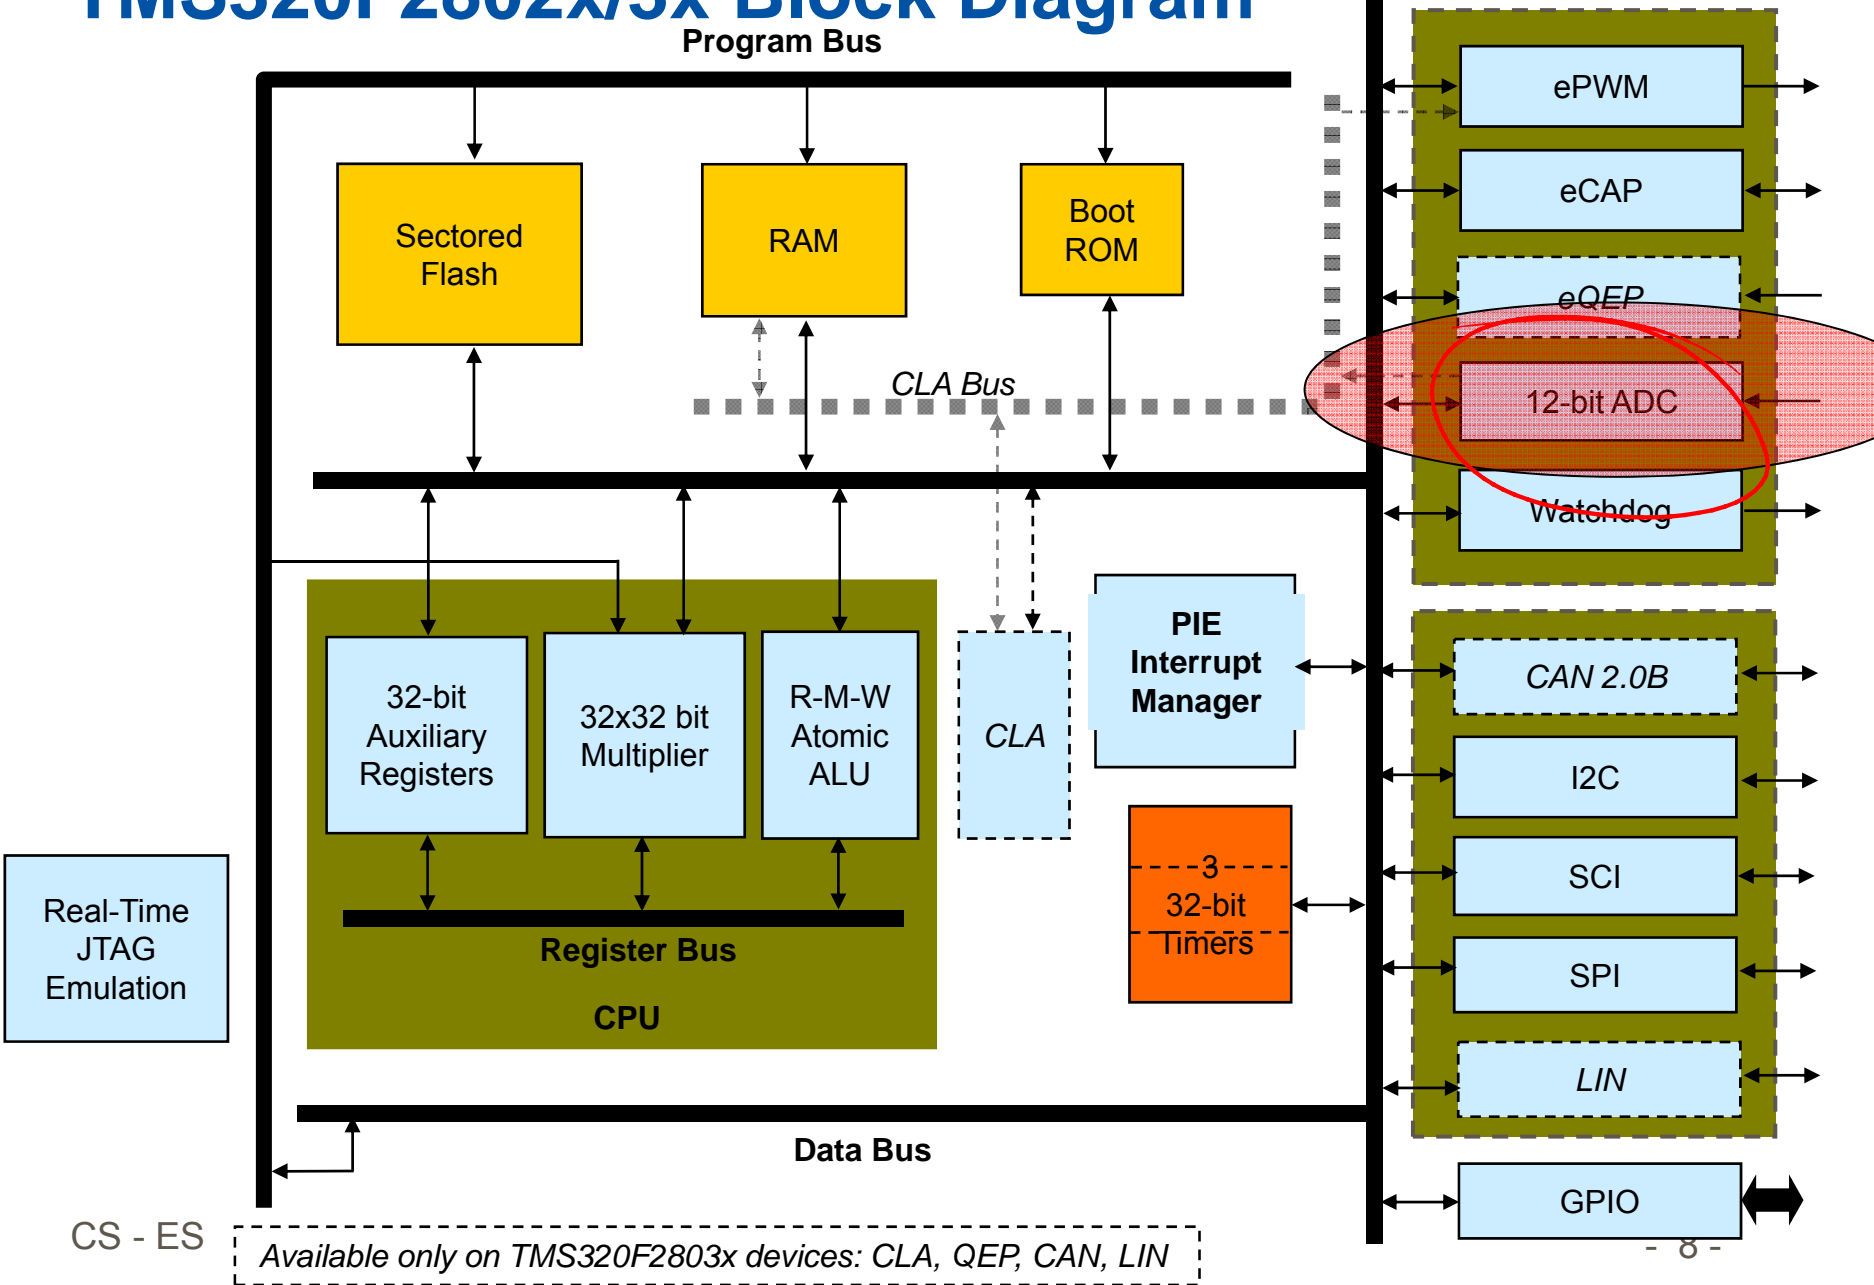


LED Lighting

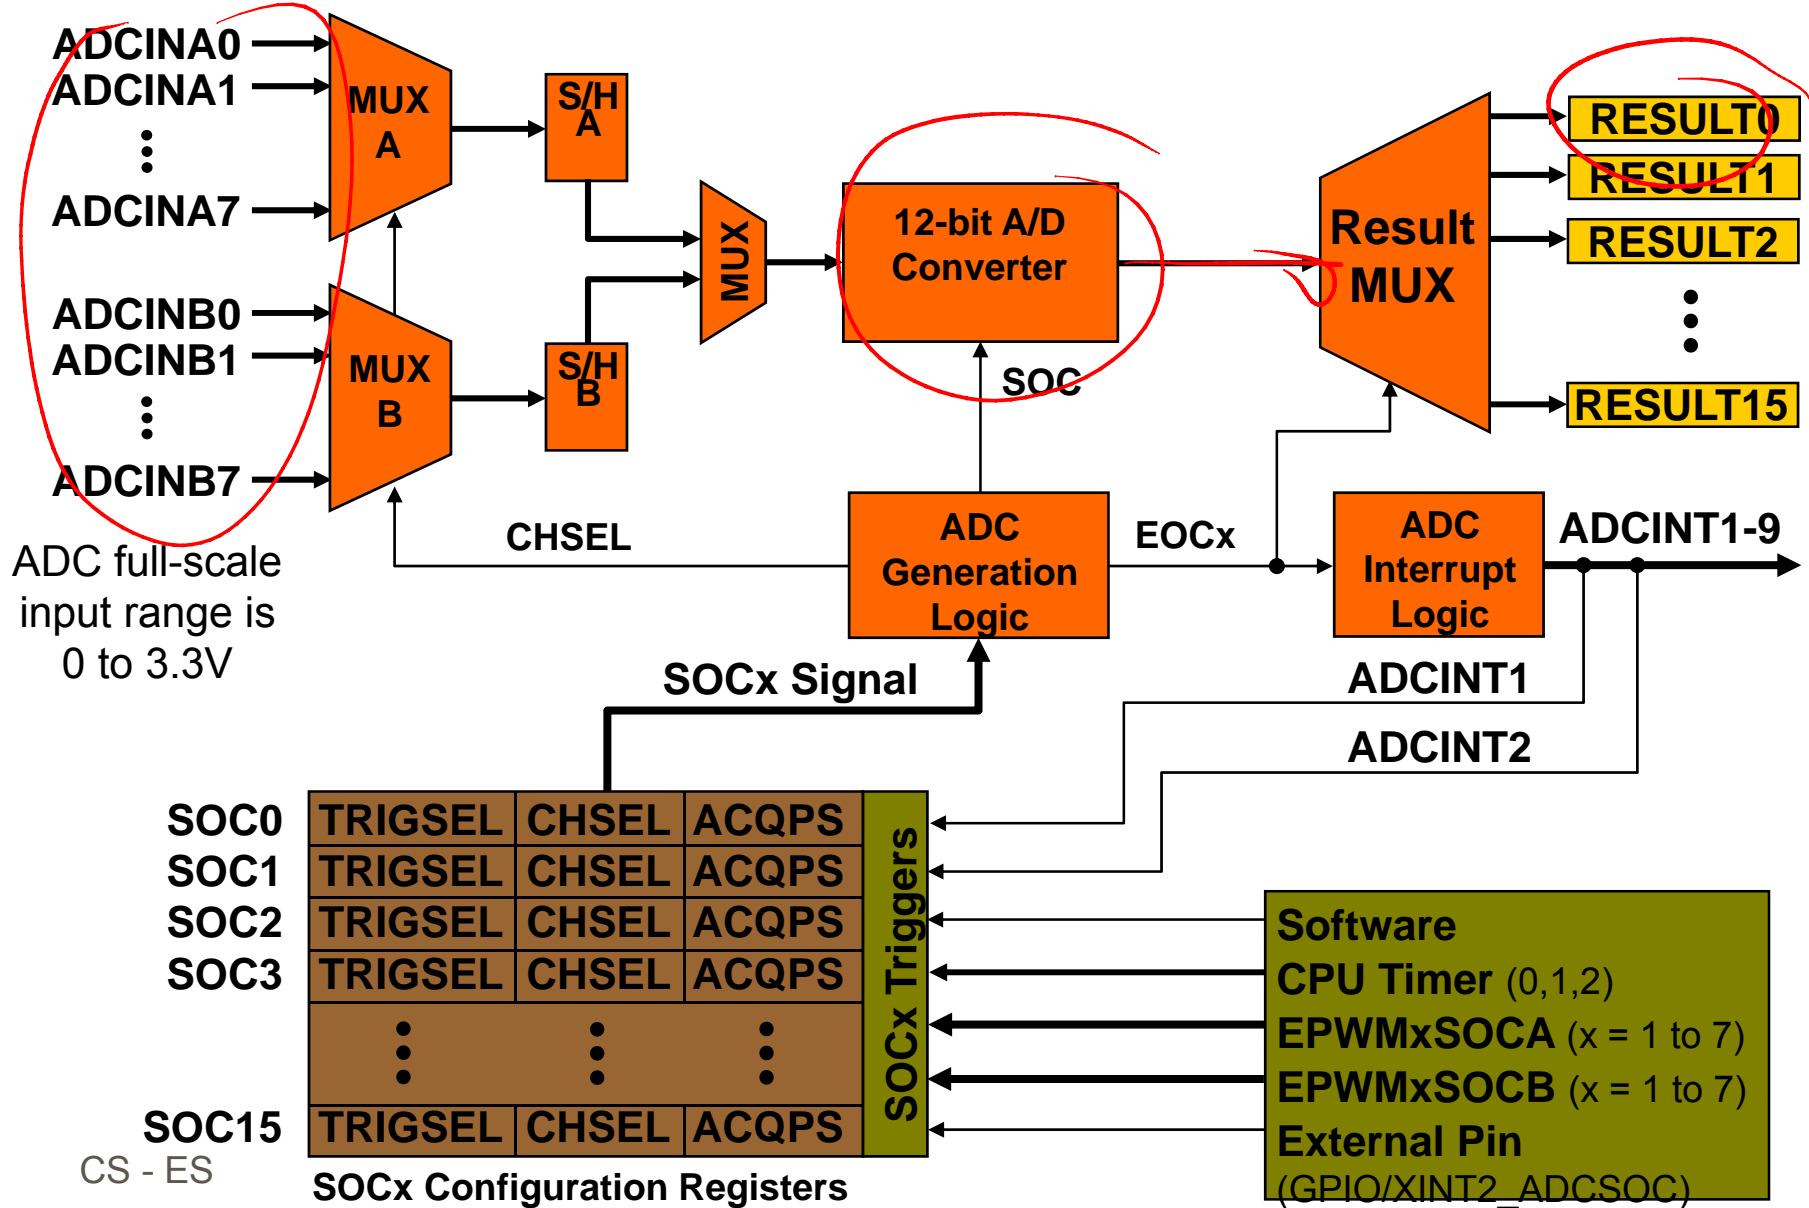


CS - ES

TMS320F2802x/3x Block Diagram

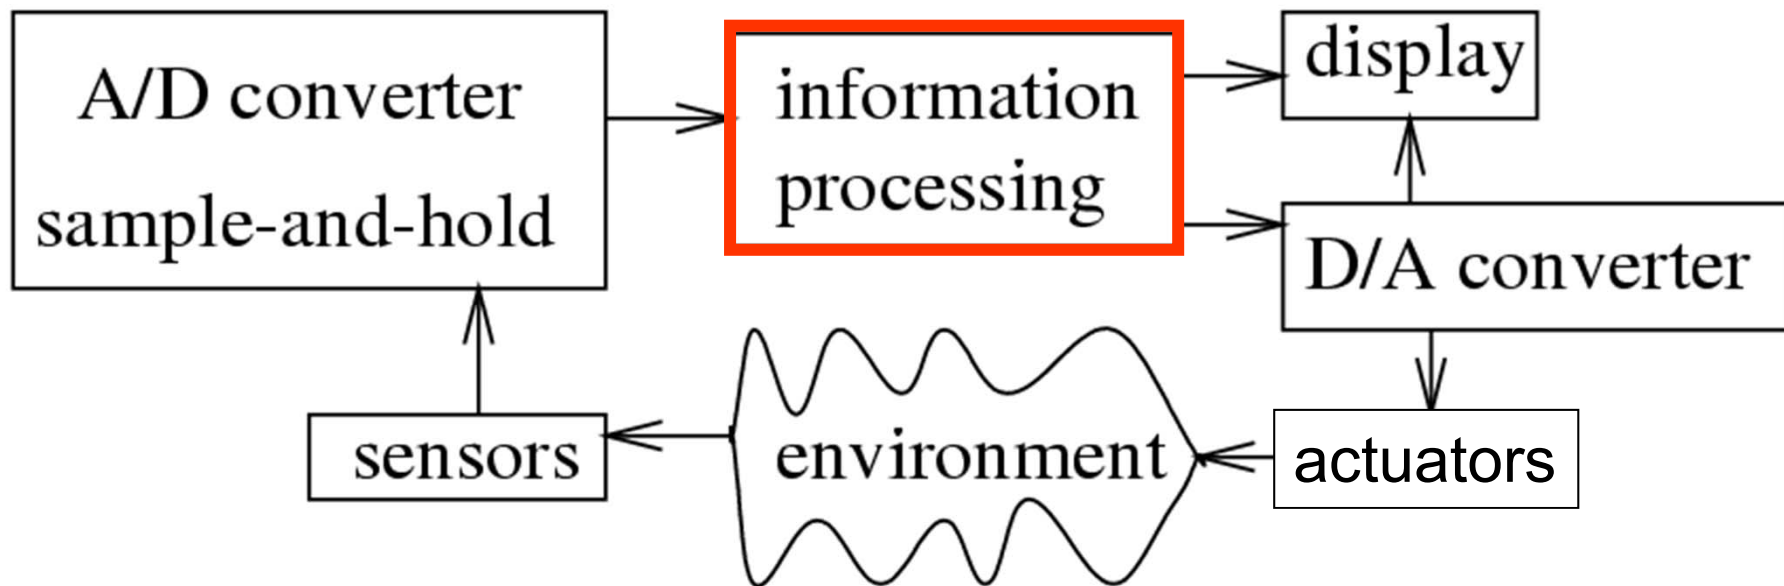


ADC Module Block Diagram



Embedded System Hardware

- Embedded system hardware is frequently used in a loop („*hardware in a loop*“):



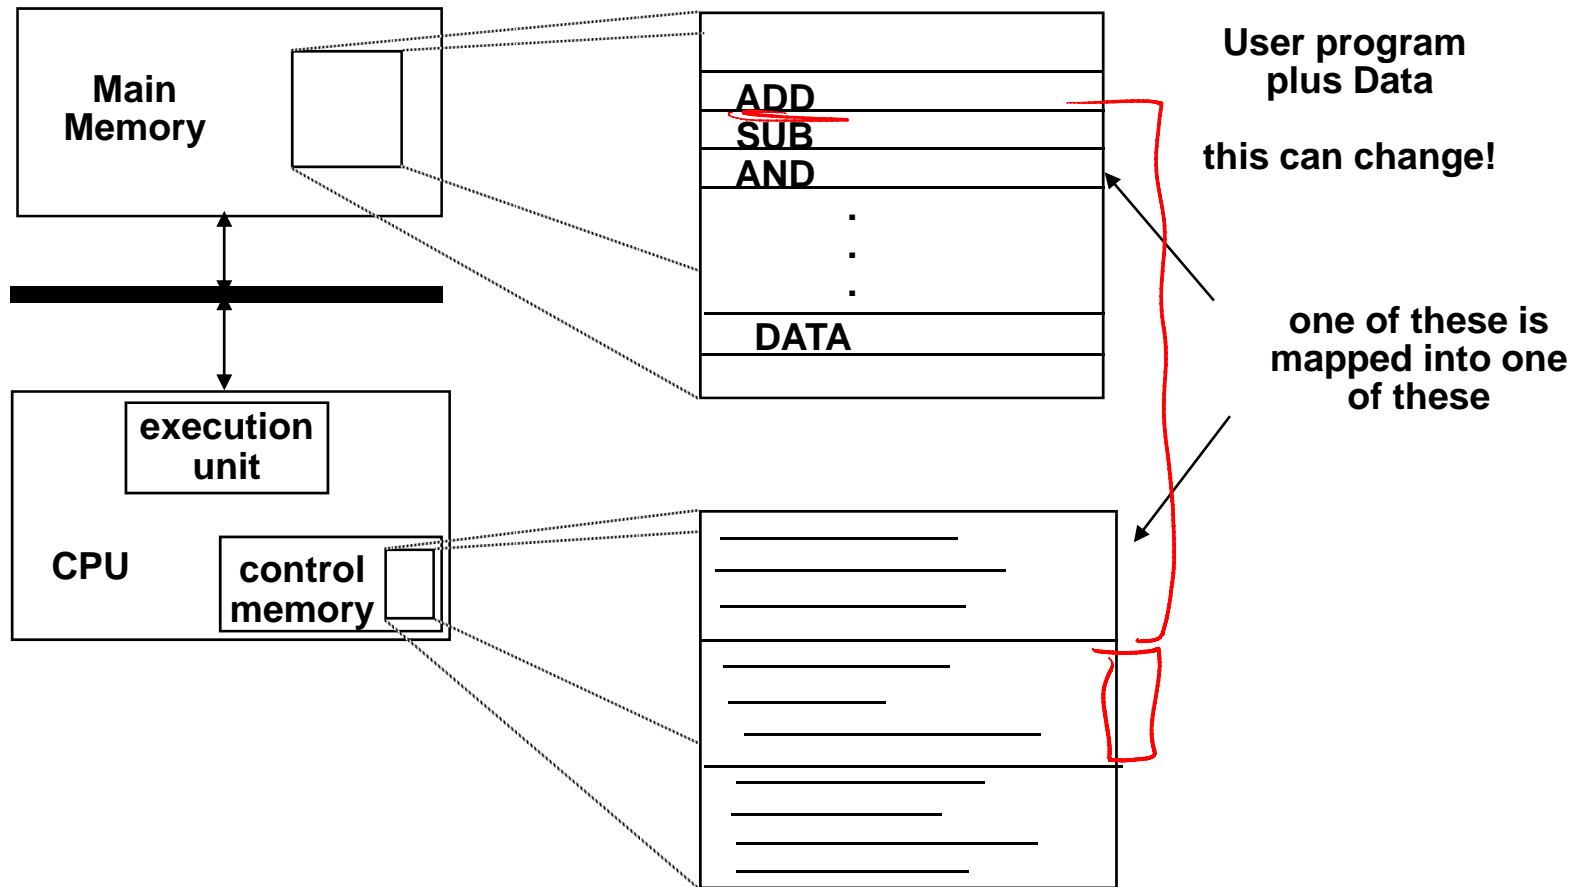
CISC vs. RISC

REVIEW

At the time of their initial development, CISC machines used available technologies to optimize computer performance.

- Microprogramming is as easy as assembly language to implement, and much less expensive than hardwiring a control unit.
- The ease of microcoding new instructions allowed designers to make **CISC machines upwardly compatible**: a new computer could run the same programs as earlier computers because the new computer would contain a superset of the instructions of the earlier computers.
- Because microprogram instruction sets can be written to **match the constructs of high-level languages**, the compiler does not have to be as complicated.

Microprogramming



Supported complex instructions a sequence of simple micro-inst

What is RISC?

- RISC, or *Reduced Instruction Set Computer*, is a type of microprocessor architecture that utilizes a **small, highly-optimized set of instructions**, rather than a more specialized set of instructions often found in other types of architectures.
- **About 80% of the computations of a typical program required only about 20% of the instructions in a processor's instruction set.** The most frequently used instructions were simple instructions such as load, store and add.
- Certain design features have been characteristic of most RISC processors:
 - **one cycle execution time**: RISC processors have a CPI (clock per instruction) of one cycle. This is due to the optimization of each instruction on the CPU and a technique called PIPELINING
 - **pipelining: a technique that allows for simultaneous execution of parts, or stages, of instructions to more efficiently process instructions;**
 - **large number of registers**: the RISC design philosophy generally incorporates a larger number of registers to prevent in large amounts of interactions with memory

RISC's disadvantages

- **Code Quality**

The performance of a RISC processor depends greatly on the code that it is executing.

If the programmer (or compiler) does a poor job of instruction scheduling, the processor can spend quite a bit stalling: waiting for the result of one instruction before it can proceed with a subsequent instruction.

Since the scheduling rules can be complicated, most programmers use a high level language (such as C or C++) and leave the instruction scheduling to the compiler.

This makes the performance of a RISC application depend critically on the quality of the code generated by the compiler. Therefore, developers (and development tool suppliers such as Apple) have to choose their compiler carefully based on the quality of the generated code.

Comparision

Feature	RISC	CISC
Power	One or two mill watts	Many watts
Compute Speed	Up to a mega-flop	Up to several mega-flop
I/O	Custom, any sort of hardware	PC based options via a BIOS
Cost	<u>Dollars</u>	Tens to hundreds of <u>Dollars</u>
Environmental	High Temp, Low EM Emissions	Needs Fans
Operating System Port	<u>Difficult</u> - Roughly equivalent to making a Mac OS run on a SPARC Station	Load and Go- simplified by an industry standard BIOS

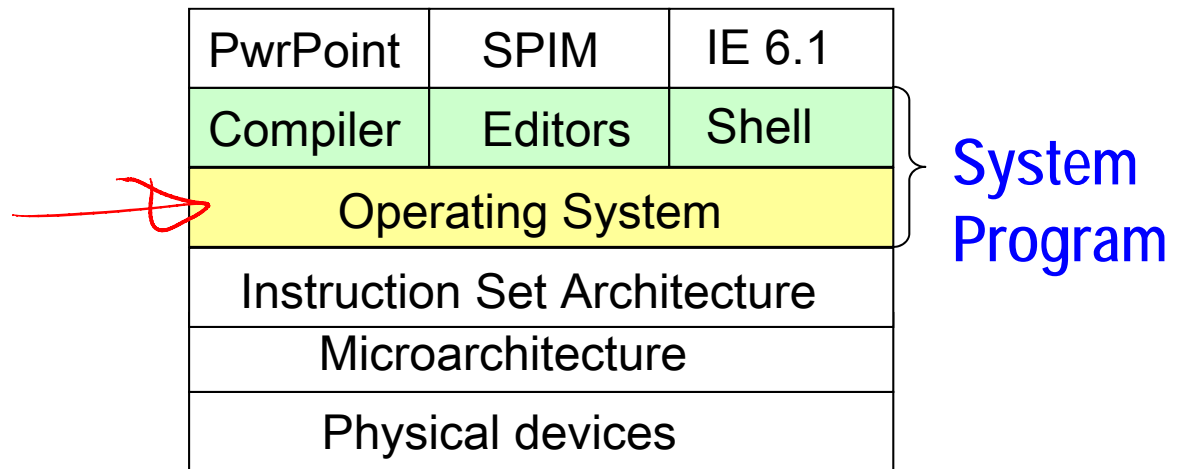
“Iron Law” of Processor Performance

$$\frac{\text{Time}}{\text{Program}} = \frac{\text{Instructions}}{\text{Program}} * \frac{\text{Cycles}}{\text{Instruction}} * \frac{\text{Time}}{\text{Cycle}}$$

- Instructions per program depends on source code, compiler technology, and ISA
- Cycles per instructions (CPI) depends upon the ISA and the microarchitecture
- Time per cycle depends upon the microarchitecture and the base technology
- RISC systems shorten execution time by **reducing the clock cycles per instruction.**
- CISC systems improve performance **by reducing the number of instructions per program.**

What is an Operating System?

- An intermediate program between a user of a computer and the computer hardware (to hide messy details)
- Goals:
 - Execute user programs and make solving user problems easier
 - Make the computer system convenient and efficient to use



Operating System Concepts

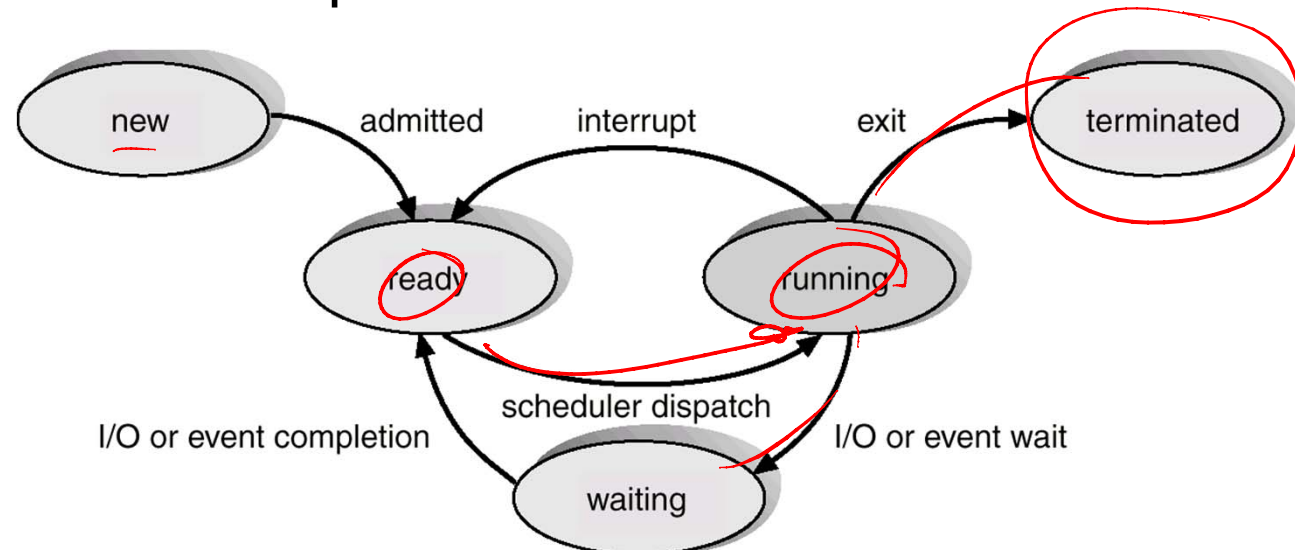
- **Process Management**
- Main Memory Management
- File Management
- I/O System Management
- Secondary Management
- Networking
- Protection System
- Command-Interpreter System

Process Management

- A *process* is a program in execution
- A process contains
 - Address space (e.g. read-only code, global data, heap, stack, etc)
 - PC, \$sp
 - Opened file handles
- A process needs certain resources, including CPU time, memory, files, and I/O devices
- The OS is responsible for the following activities for process management
 - Process creation and deletion
 - Process suspension and resumption
 - Provision of mechanisms for:
 - process synchronization
 - process communication

Process State

- As a process executes, it changes *state*
 - new: The process is being created
 - ready: The process is waiting to be assigned to a process
 - running: Instructions are being executed
 - waiting: The process is waiting for some event (e.g. I/O) to occur
 - terminated: The process has finished execution

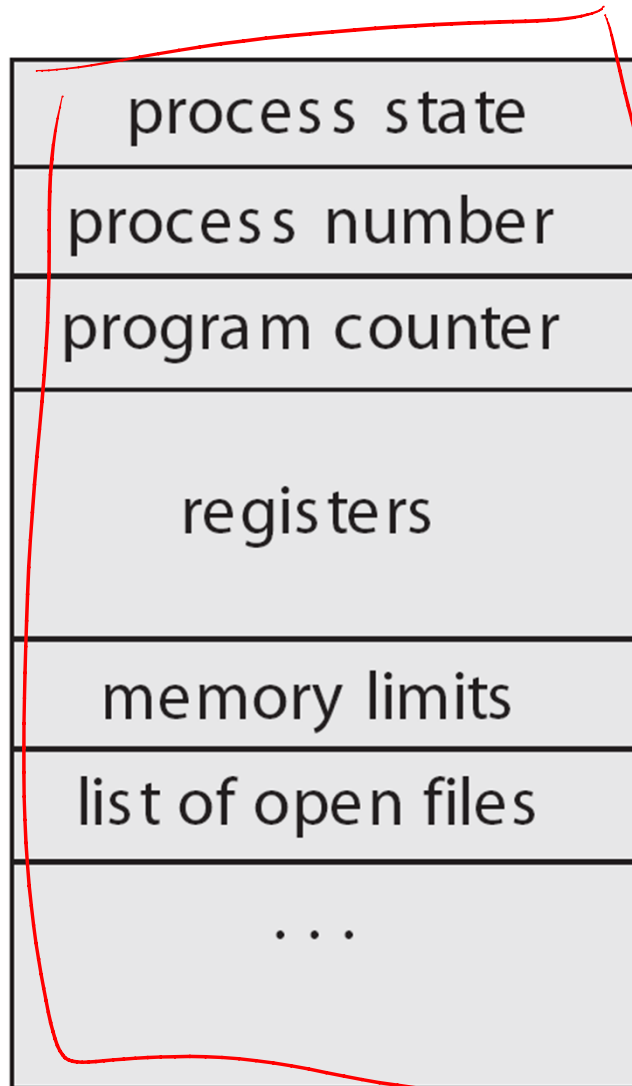


Process Control Block (PCB)

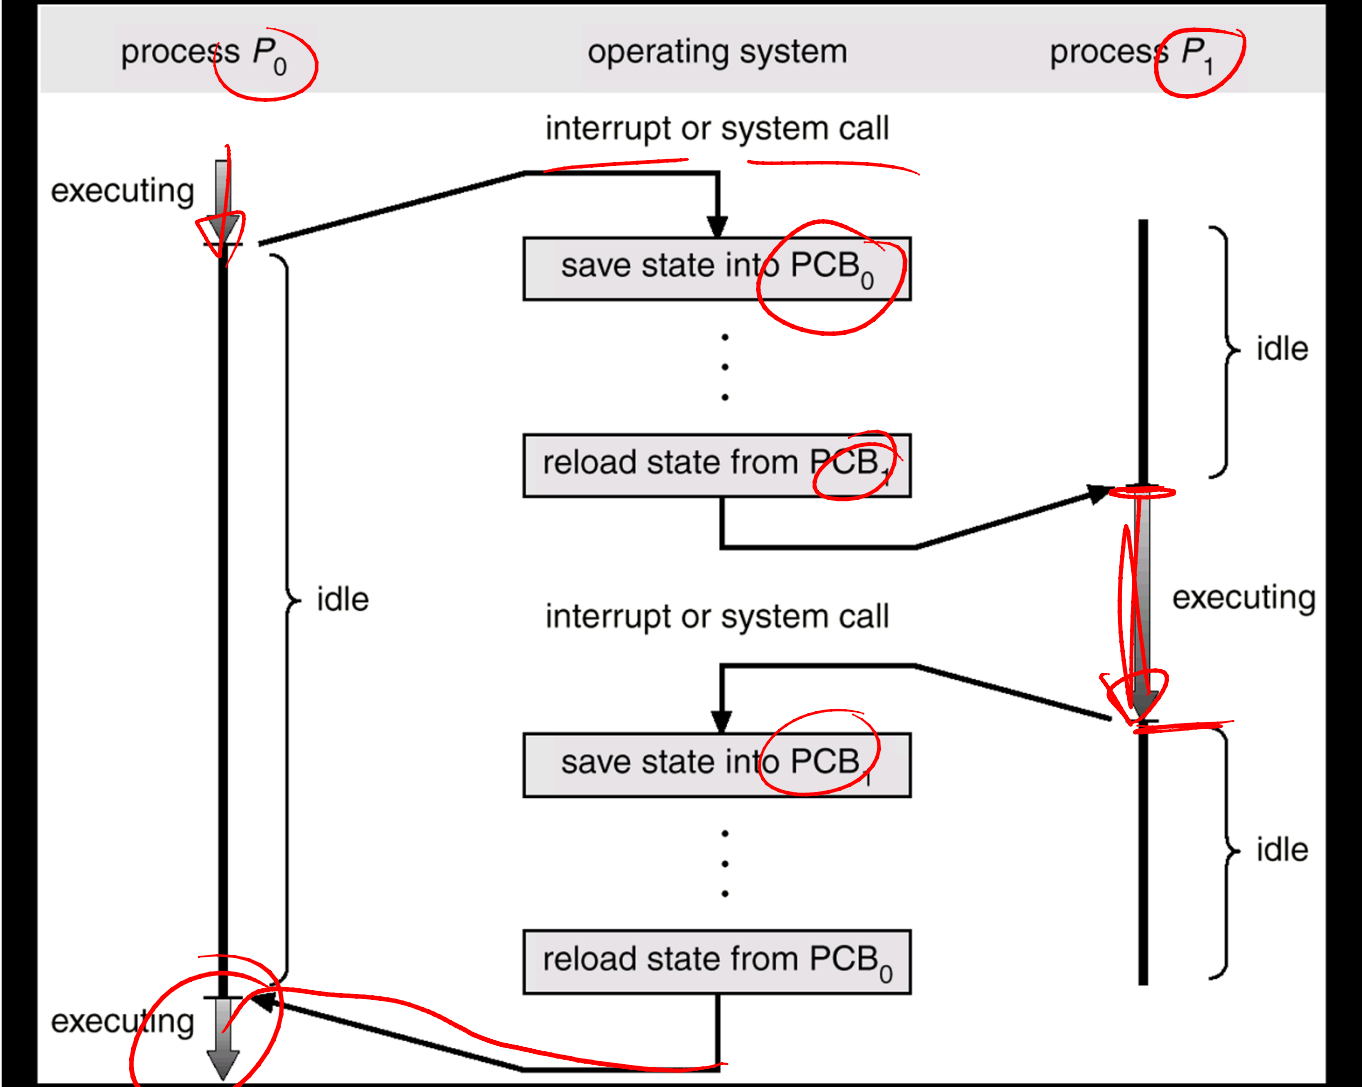
Information associated with each process

- Process state
- Program counter
- CPU registers (for context switch)
- CPU scheduling information (e.g. priority)
- Memory-management information (e.g. page table, segment table)
- Accounting information (PID, user time, constraint)
- I/O status information (list of I/O devices allocated, list of open files etc.)

Process Control Block (PCB)



CPU Switch From Process to Process

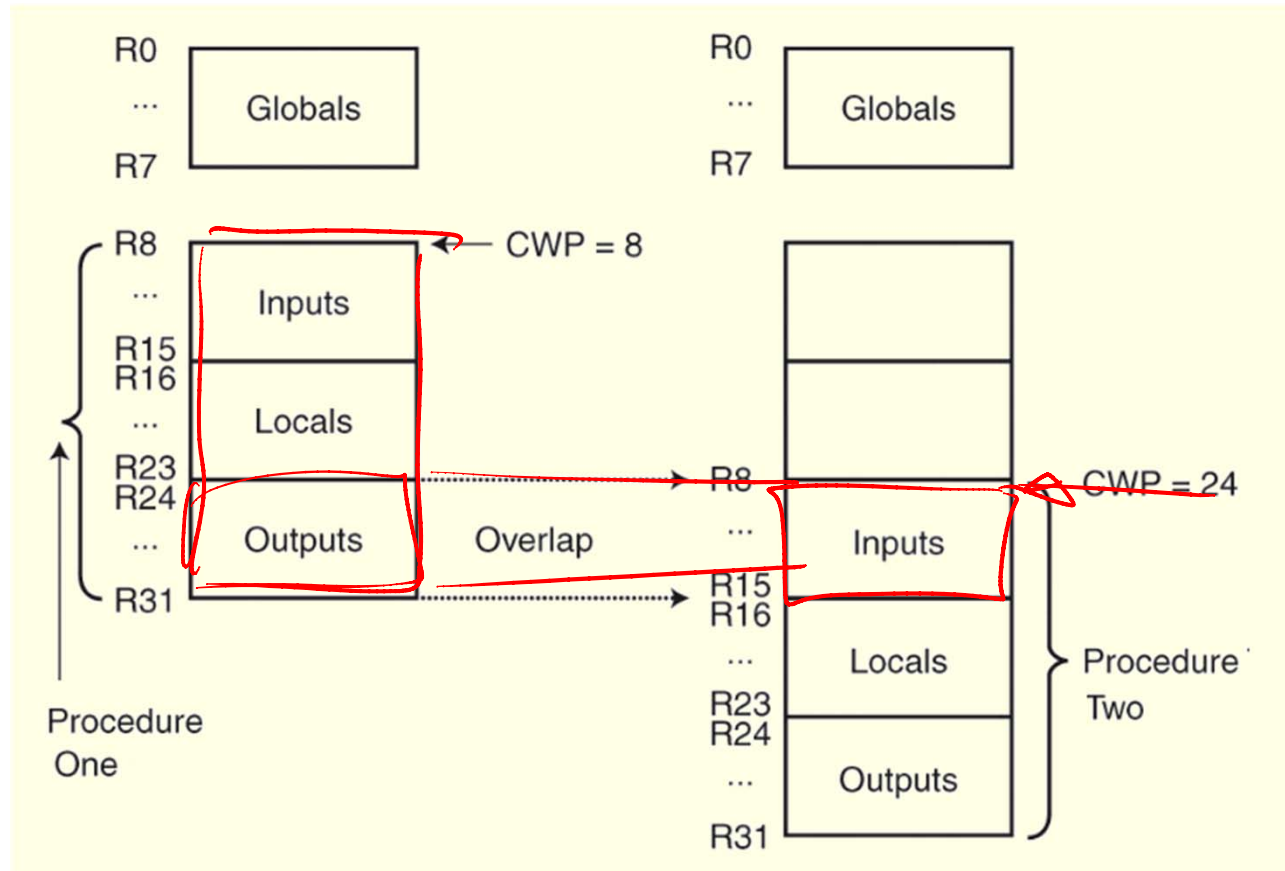


RISC Machines

- Because of their load-store ISAs, RISC architectures require a **large number of CPU registers.**
- These registers provide **fast access to data during sequential program execution.**
- They can also be employed to **reduce the overhead typically caused by passing parameters to subprograms.**
- Instead of pulling parameters off of a **stack**, the subprogram is directed to use a **subset of registers.**
- **Fast Context Switching** - support with **two additional local register banks** (e.g.; Infineon **XC167CI**)
- E.g.; Berkeley RISC: **> 100 Regs**
only 32 visible for the program.

RISC Machines

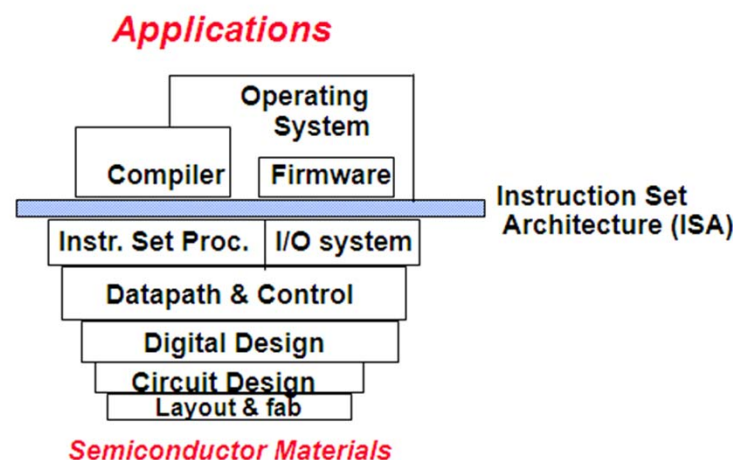
- This is how registers can be overlapped in a RISC system.
- The *current window pointer (CWP)* points to the active register window.



Instruction Set Architecture

Is the interface between hardware and software.

- allows easy programming (compilers, OS, ..);
 - Provides convenient functionality to higher levels
- allows efficient implementations (hardware);
 - Permits an efficient implementation at lower levels
- has a long lifetime (survives many HW generations) - portability



Instruction Set Architecture (ISA) versus Implementation

- ISA is the hardware/software interface
 - Defines set of programmer visible state
 - Defines instruction format (bit encoding) and instruction semantics
 - Examples: *MIPS, x86, IBM 360, JVM*
- Many possible implementations of one ISA
 - 360 implementations: model 30 (c. 1964), z990 (c. 2004)
 - x86 implementations: *8086 (c. 1978), 80186, 286, 386, 486, Pentium, Pentium Pro, Pentium-4 (c. 2000), AMD Athlon, Transmeta Crusoe, SoftPC*
 - MIPS implementations: *R2000, R4000, R10000, ...*
 - JVM: *HotSpot, PicoJava, ARM Jazelle, ...*

Styles of ISA

- Accumulator
- Stack
- GPR

- CISC
- RISC
- VLIW
- Vector

- Boundaries are fuzzy, and hybrids are common
 - E.g., 8086/87 is hybrid accumulator-GPR-stack ISA
 - Many ISAs have added vector extensions



**XC167
Derivatives**

Preliminary

Functional Description

3 Functional Description

The architecture of the XC167 combines advantages of RISC, CISC, and DSP processors with an advanced peripheral subsystem in a very well-balanced way. In addition, the on-chip memory blocks allow the design of compact systems-on-silicon with maximum performance (computing, control, communication).

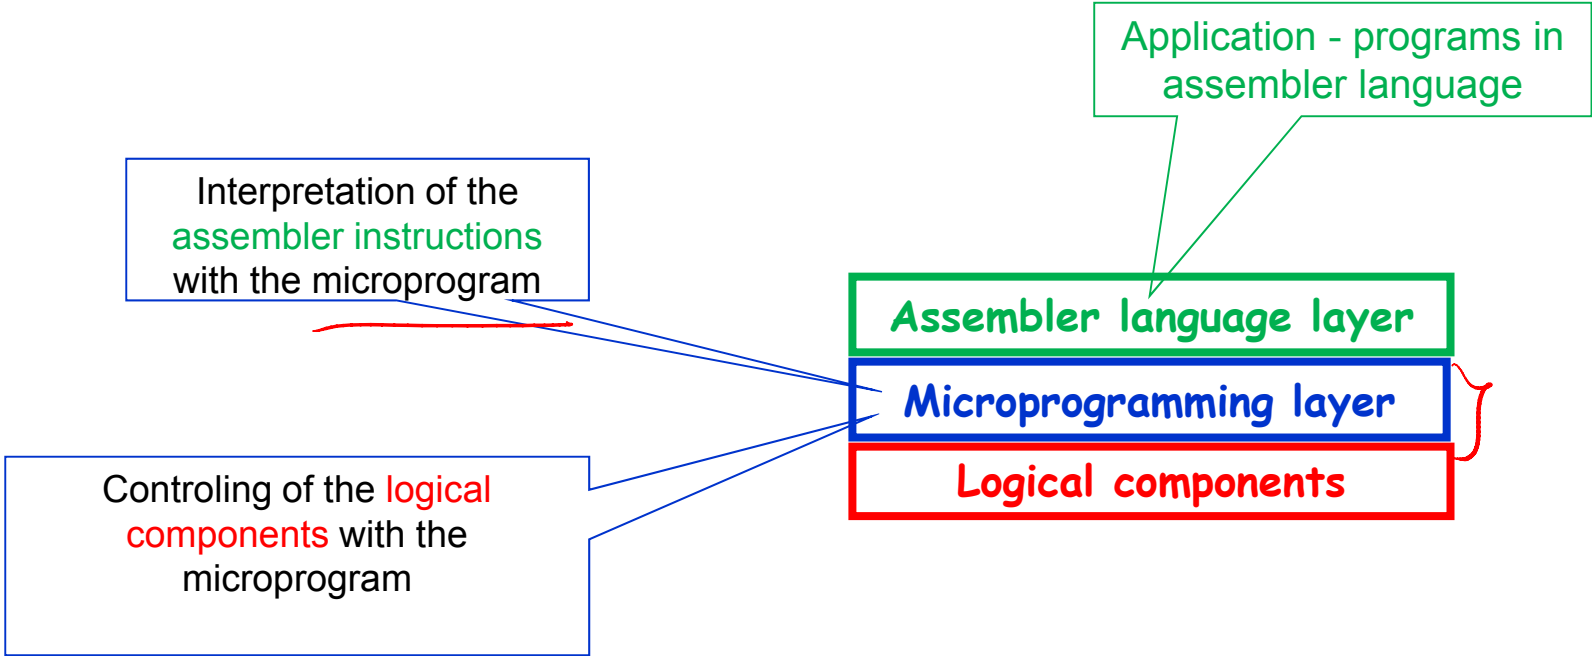
Styles of Implementation

- Microcoded

- Unpipelined single cycle
- Hardwired in-order pipeline
- Software interpreter
- Just-in-Time compiler

Micro programming

Tasks of the MP layer



Microarchitecture

Registerbank

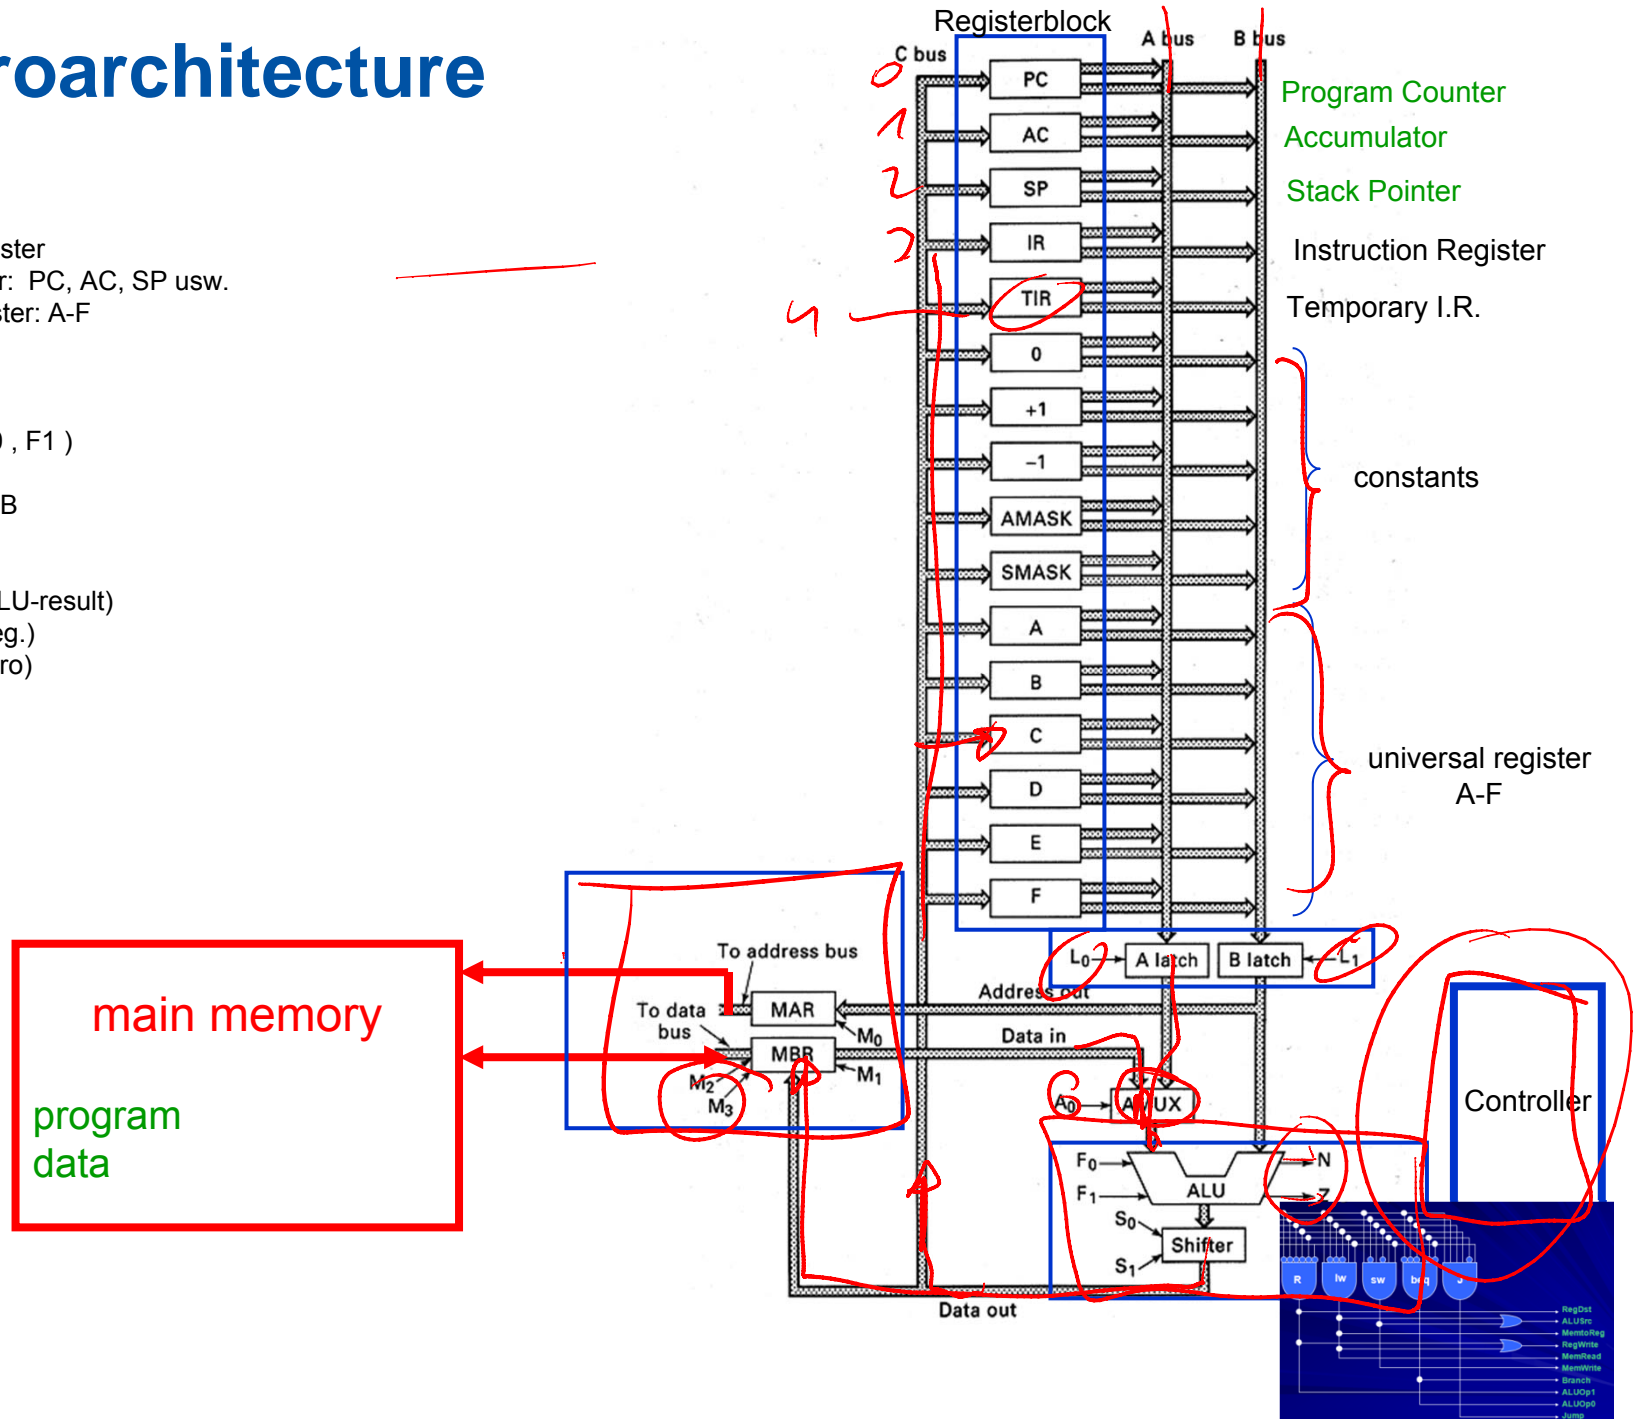
16 16-Bit register
 special register: PC, AC, SP usw.
 universalRegister: A-F

ALU

16-Bit
 4 funktions (F0, F1)
 A + B
 A and B
 \bar{A}
 A
 2 statusbits (ALU-result)
 N (neg.)
 Z (zero)

shifter

1 Bit t



Format micro instruction

Signals for data path and memory:

16	control signals	load A-Bus
16	control signals	load B-Bus
16	"-	load C-Bus
2	"-	A, B- Latch
2	"-	ALU-functions
2	"-	shifter
1	"-	MAR (M0)
3	"-	<u>MBR (M1)</u> , memory read/write (M2, M3)
1	"-	<u>AMUX (A0)</u>
1	"-	Enable C-Bus (ENC)

60 Bit per micro instruction




Format micro instruction

Reduction of the number of control bits

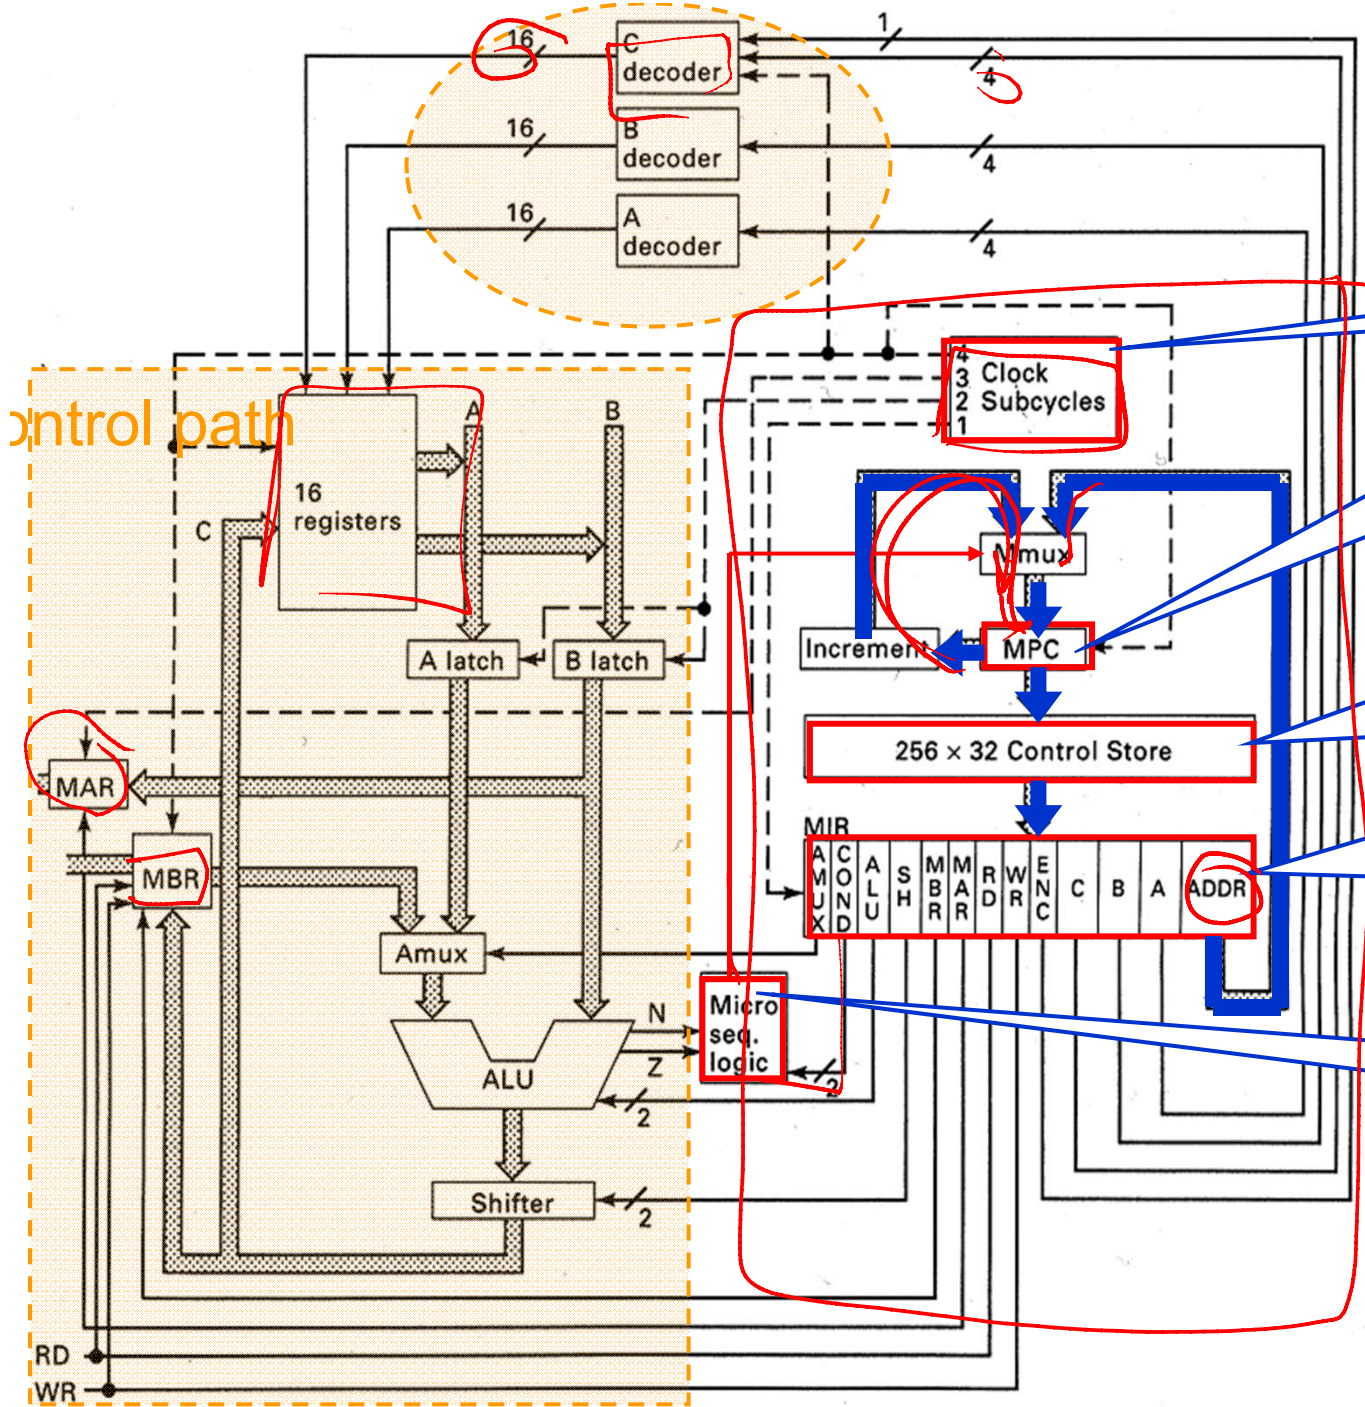
Use coding

A-Bus 4 Bit (instead of 16)
 B-Bus 4 Bit
 C-Bus 4 Bit



- AMUX —controls left ALU input: 0 = A latch, 1 = MBR
- ALU —ALU function: 0 = A + B, 1 = A AND B, 2 = A, 3 = \bar{A}
- SH —shifter function: 0 = no shift, 1 = right, 2 = left
- MBR —load MBR from shifter: 0 = don't load MBR, 1 = load MBR
- MAR —load MAR from B latch: 0 = don't load MAR, 1 = load MAR
- RD —requests memory read: 0 = no read, 1 = load MBR from memory
- WR —requests memory write: 0 = no write, 1 = write MBR to memory
- ENC —controls storing into scratchpad: 0 = don't store, 1 = store
- C —selects register for storing into if ENC = 1: 0 = PC, 1 = AC, etc.
- B —selects B bus source: 0 = PC, 1 = AC, etc.
- A —selects A bus source: 0 = PC, 1 = AC, etc.

⇒ control unit



control path

Microprogram control unit

Clock generator (4-phases)

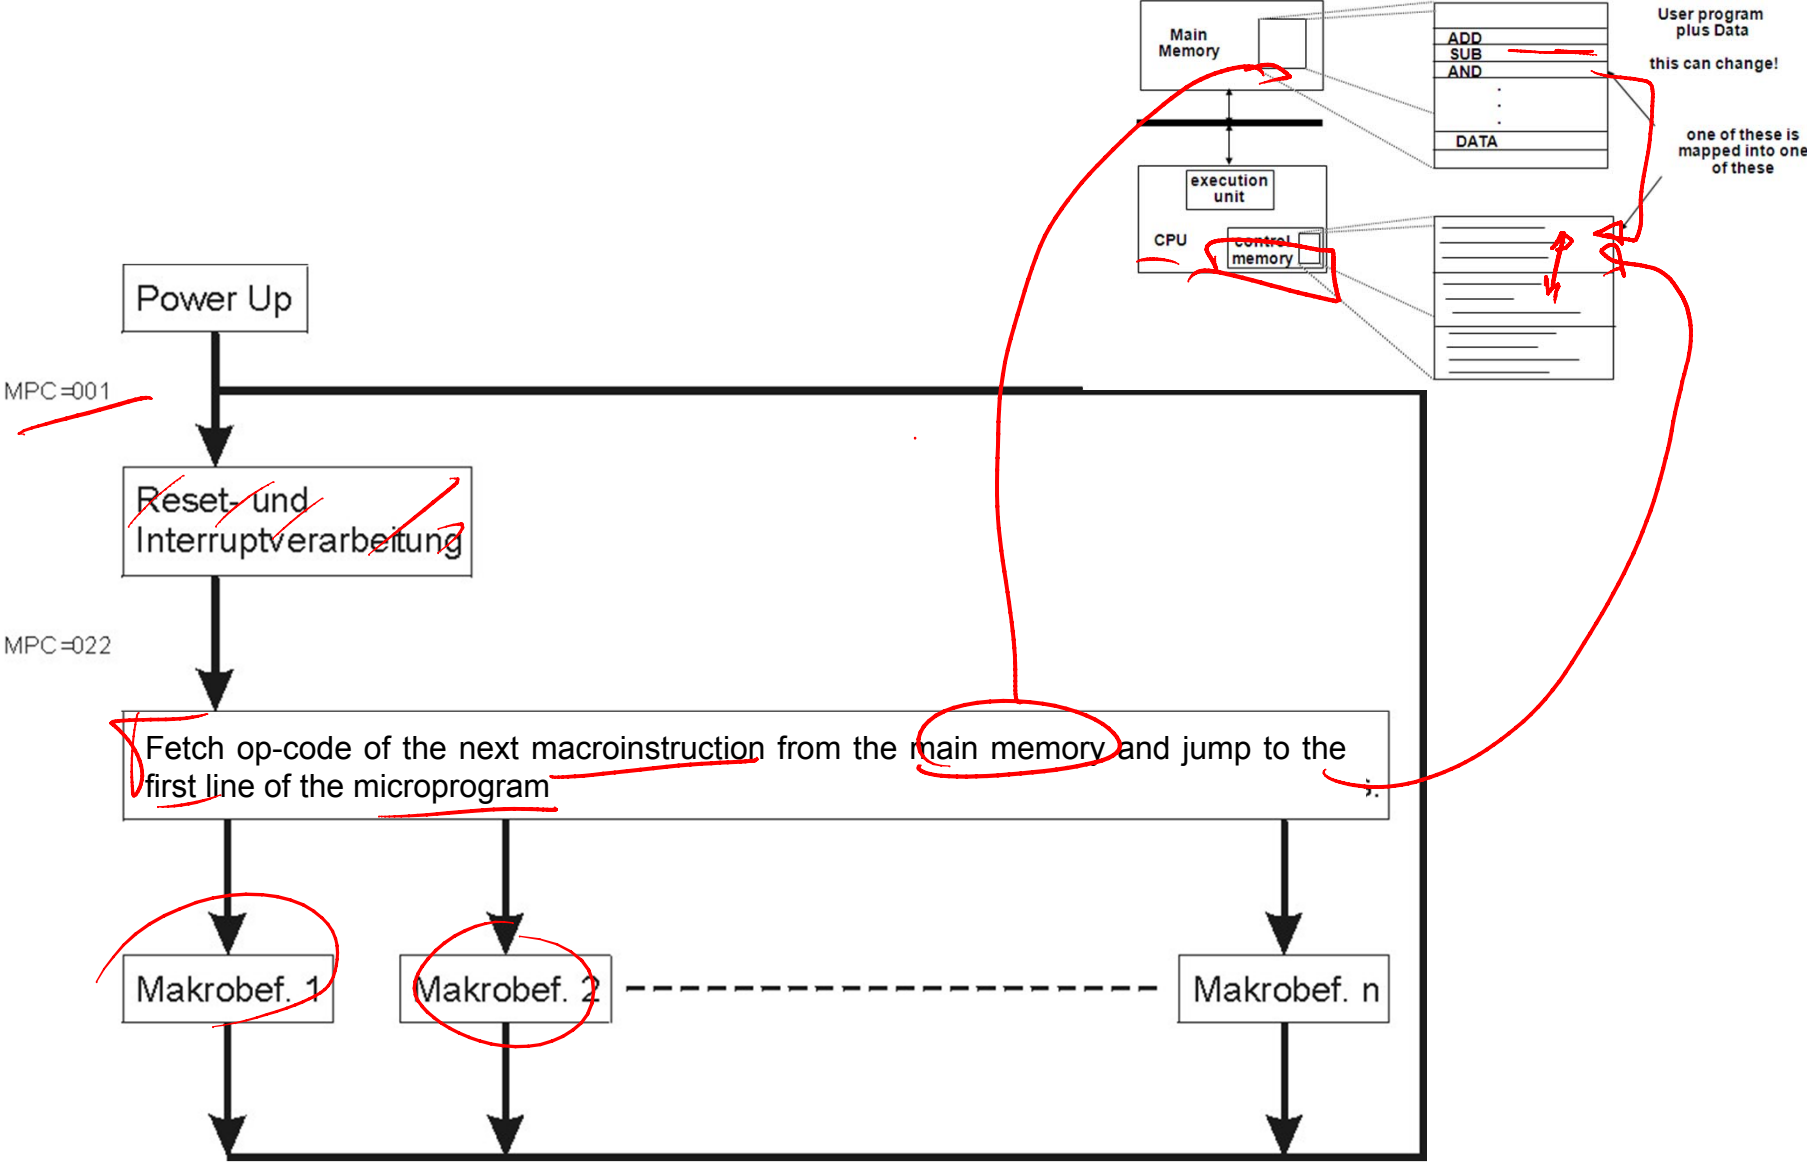
(MPC)

Microprogram memory
(256 words x 32 Bits):
Stores the micro program

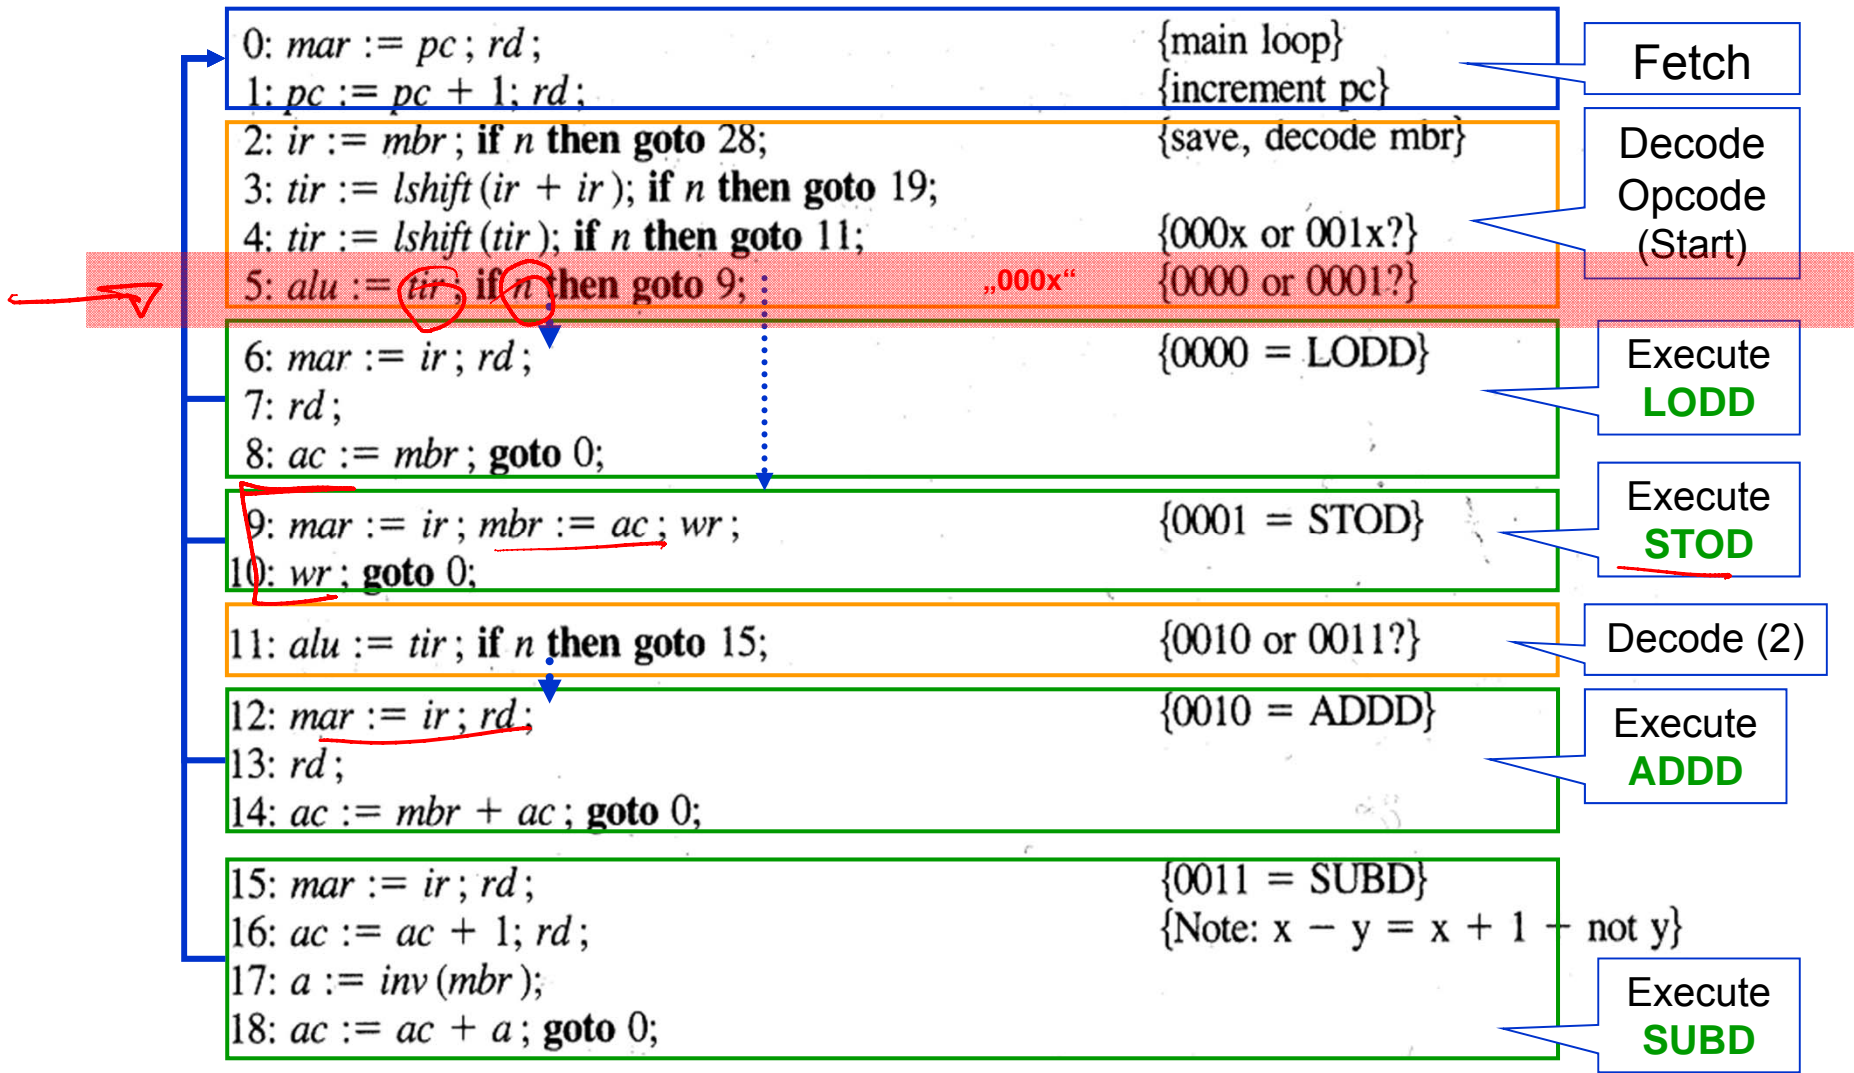
Microinstruction register (MIR)

Microsequencer
(„next-Address-Logic“)

Interpretation – macroinstruction



Microprogramm ("Interpreter") for the macroarchitecture




microinstruction

Register-Transfer-Notation

	A	C	A	S	M	M	R	W	E	C	B	A	ADDR
	U	N	L	H	B	A	R	D	N				
	X	D	U		R	R	D	D	C				
<i>mar := pc; rd;</i>	0	0	2	0	0	1	1	0	0	0	0	0	00
<i>rd;</i>	0	0	2	0	0	0	1	0	0	0	0	0	00
<i>ir := mbr</i>	1	0	2	0	0	0	0	0	1	3	0	0	00
<i>pc := pc + 1</i>	0	0	0	0	0	0	0	0	1	0	6	0	00
<i>mar := ir; mbr := ac; wr;</i>	0	0	2	0	1	1	0	1	0	0	3	1	00
<i>alu := tir; if n then goto 15;</i>	0	1	2	0	0	0	0	0	0	0	0	4	15
<i>ac := inv (mbr);</i>	1	0	3	0	0	0	0	0	1	1	0	0	00
<i>tir := l shift (tir); if n then goto 25;</i>	0	1	2	2	0	0	0	0	1	4	0	4	25
<i>alu := ac; if z then goto 22;</i>	0	2	2	0	0	0	0	0	0	0	0	1	22
<i>ac := band (ir, amask); goto 0</i>	0	3	1	0	0	0	0	0	1	1	8	3	00
<i>sp := sp + (-1); rd;</i>	0	0	0	0	0	0	1	0	1	2	2	7	00
<i>tir := l shift (ir + ir); if n then goto 69</i>	0	1	0	2	0	0	0	0	1	4	3	3	69

Horizontal vs Vertical μ Code



- Horizontal μ code has wider μ instructions
 - Multiple parallel operations per μ instruction
 - Fewer steps per macroinstruction
 - Sparser encoding \Rightarrow more bits
- Vertical μ code has narrower μ instructions
 - Typically a single datapath operation per μ instruction
 - separate μ instruction for branches
 - More steps to per macroinstruction
 - More compact \Rightarrow less bits
- Nanocoding
 - Tries to combine best of horizontal and vertical μ code

Dictionary approach, two level control store (indirect addressing of instructions)

“Dictionary-based coding schemes cover a wide range of various coders and compressors.

Their common feature is that the methods use some kind of a dictionary that contains parts of the input sequence which frequently appear.

The encoded sequence in turn contains references to the dictionary elements rather than containing these over and over.”

[Á. Beszédés et al.: Survey of Code size Reduction Methods, Survey of Code-Size Reduction Methods, *ACM Computing Surveys*, Vol. 35, Sept. 2003, pp 223-267]

Nanocoding


Exploits recurring control signal patterns in μ code, e.g.,

$ALU_0 A \leftarrow Reg[rs]$

...

$ALU_{i_0} A \leftarrow Reg[rs]$

...



- MC68000 had 17-bit μ code containing either 10-bit μ jump or 9-bit nanoinstruction pointer
 - Nanoinstructions were 68 bits wide, decoded to give 196 control signals

Microprogramming in Modern Usage


- *Microprogramming is far from extinct*
- Played a crucial role in micros of the Eighties
DEC uVAX, Motorola 68K series, Intel 386 and 486
- Microcode plays an assisting role in most modern micros (*AMD Athlon, Intel Core 2 Duo, IBM PowerPC*)
 - Most instructions are executed directly, i.e., with hard-wired control
 - Infrequently-used and/or complicated instructions invoke the microcode engine
- **Patchable microcode** common for post-fabrication bug fixes, e.g. Intel Pentiums load μ code patches at bootup

Pipelining

Review: Single-cycle Processor

- Five steps to design a processor:

1. Analyze instruction set → datapath requirements
2. Select set of datapath components & establish clock methodology
3. Assemble datapath meeting the requirements
4. Analyze implementation of each instruction to determine setting of control points that effects the register transfer.
5. Assemble the control logic
 - Formulate Logic Equations
 - Design Circuits



Single Cycle Performance




- Assume time for actions are
 - 100ps for register read or write; 200ps for other events
- Clock rate is?

Instr	Instr fetch	Register read	ALU op	Memory access	Register write	Total time
lw	200ps	100 ps	200ps	200ps	100 ps	800ps
sw	200ps	100 ps	200ps	200ps		700ps
R-format	200ps	100 ps	200ps		100 ps	600ps
beq	200ps	100 ps	200ps			500ps

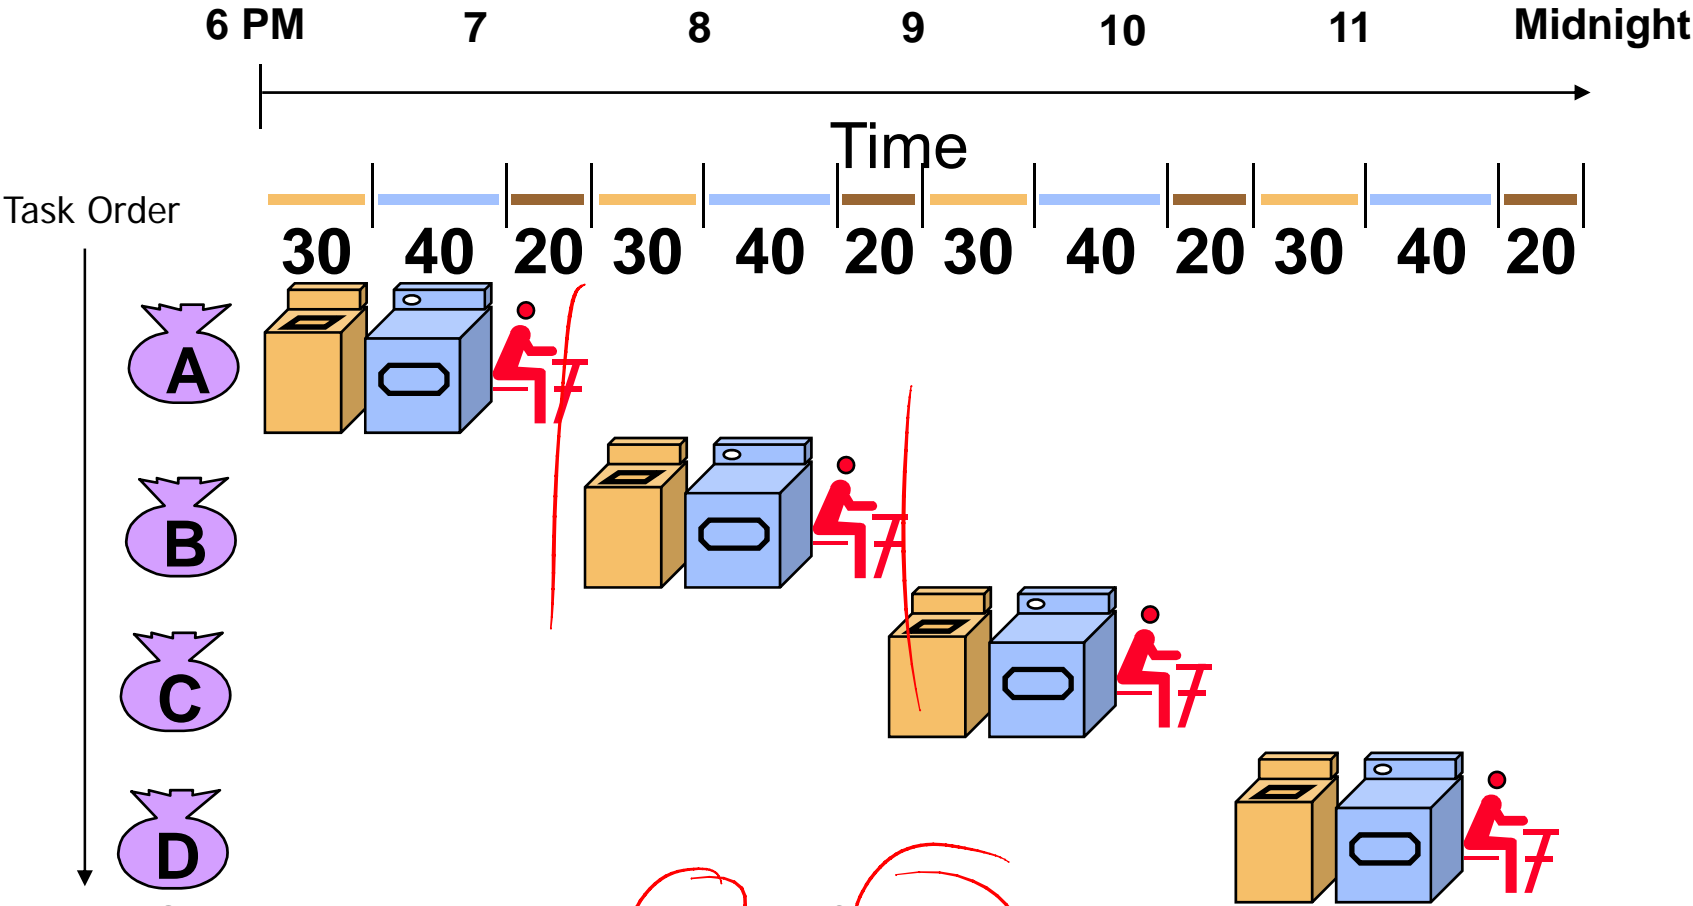
- What can we do to improve clock rate?
- Will this improve performance as well?

Pipelining: It's Natural!

- Laundry Example
 - Ann, Brian, Cathy, Dave each have one load of clothes to wash, dry, and fold
 - Washer takes 30 minutes
 - Dryer takes 40 minutes
 - “Folder” takes 20 minutes

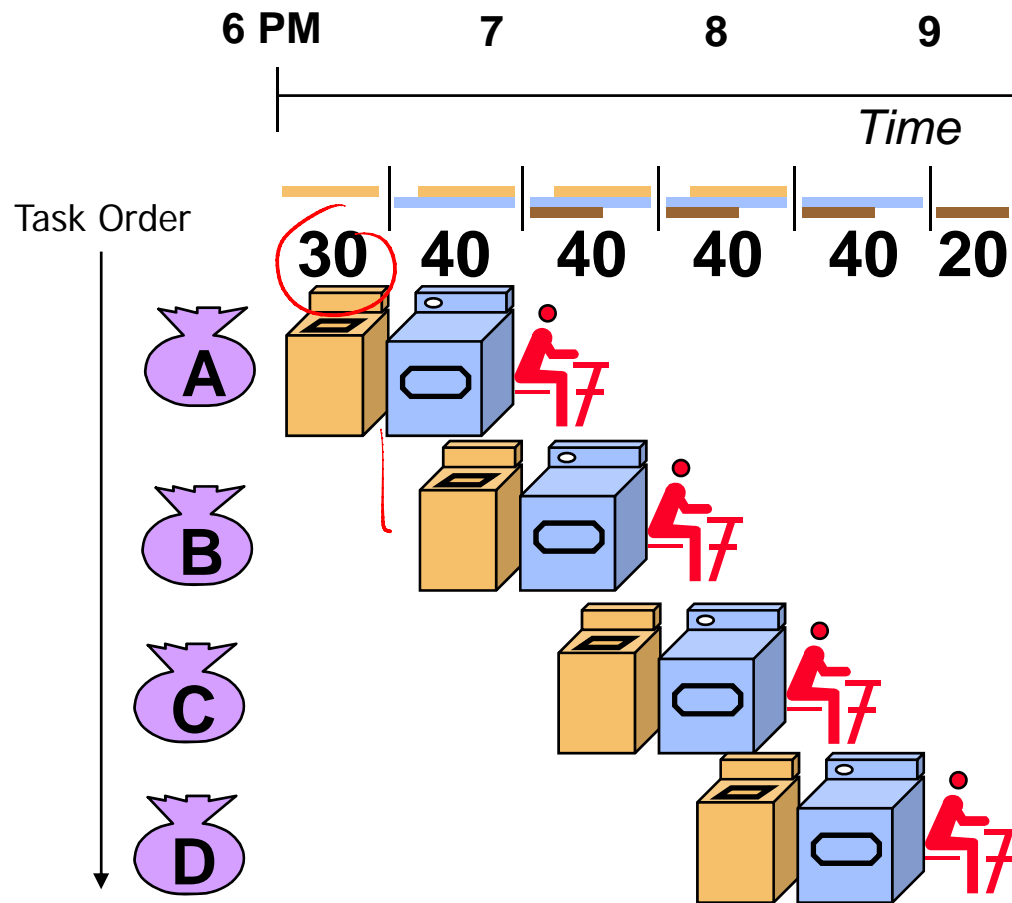


Sequential Laundry



- Sequential laundry takes 6 hours for 4 loads
- If they learned pipelining, how long would laundry take?


Pipelined Laundry: Why Wait?



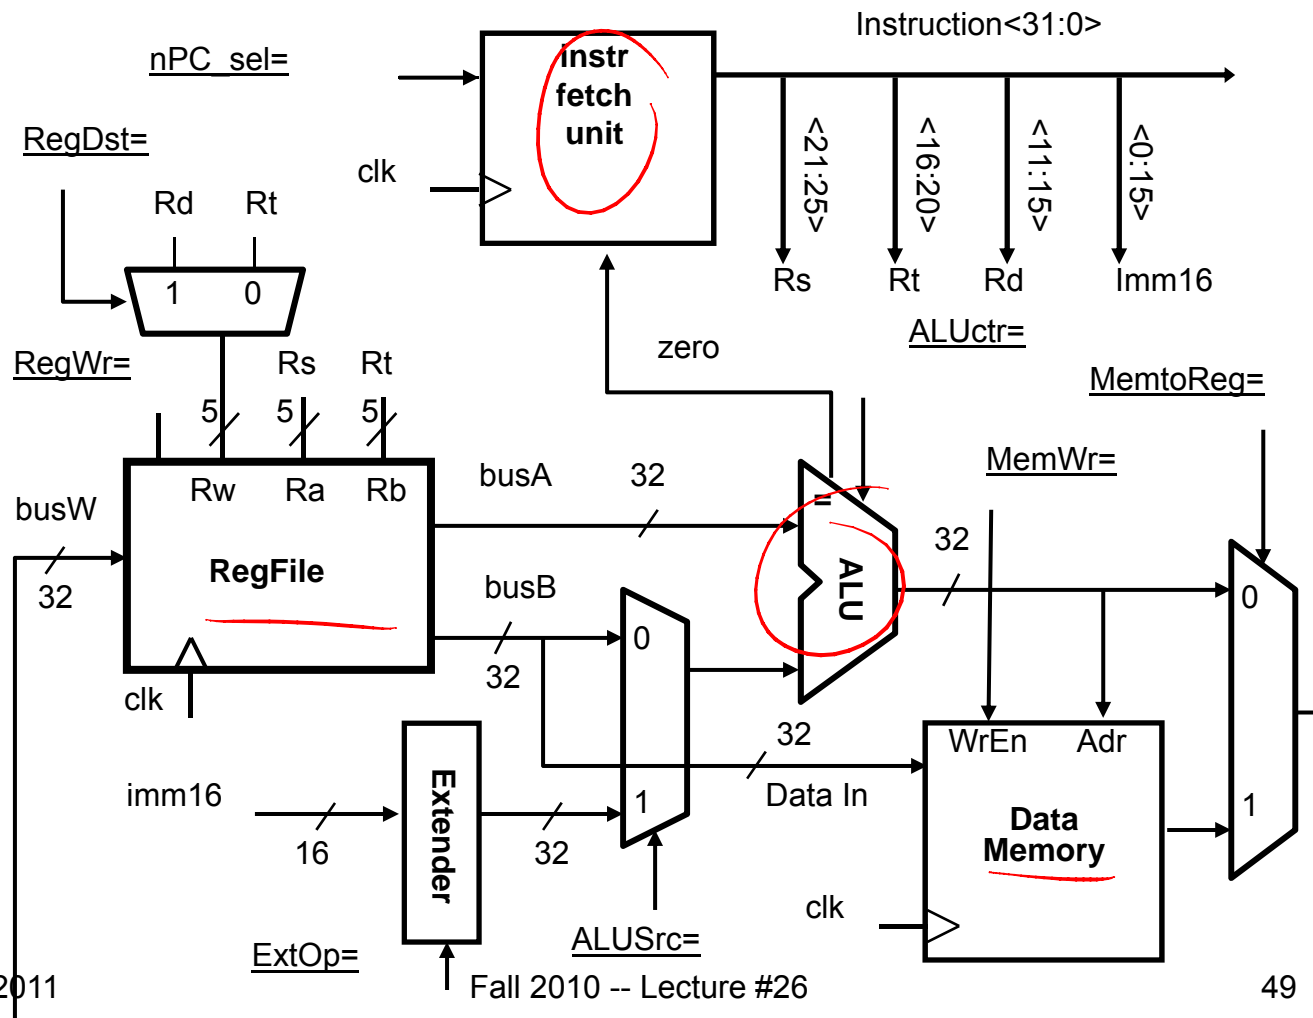
- Pipelining does **not help latency** of a single task, it **helps throughput** of entire workload
- **Multiple tasks are operating simultaneously**
- Pipeline efficiency is limited by **slowest** pipeline stage
- Potential **speedup = Number of pipeline stages**
- Unbalanced lengths of pipe stages reduces speedup

- Pipelined laundry takes 3.5 hours for 4 loads

Single Cycle Datapath



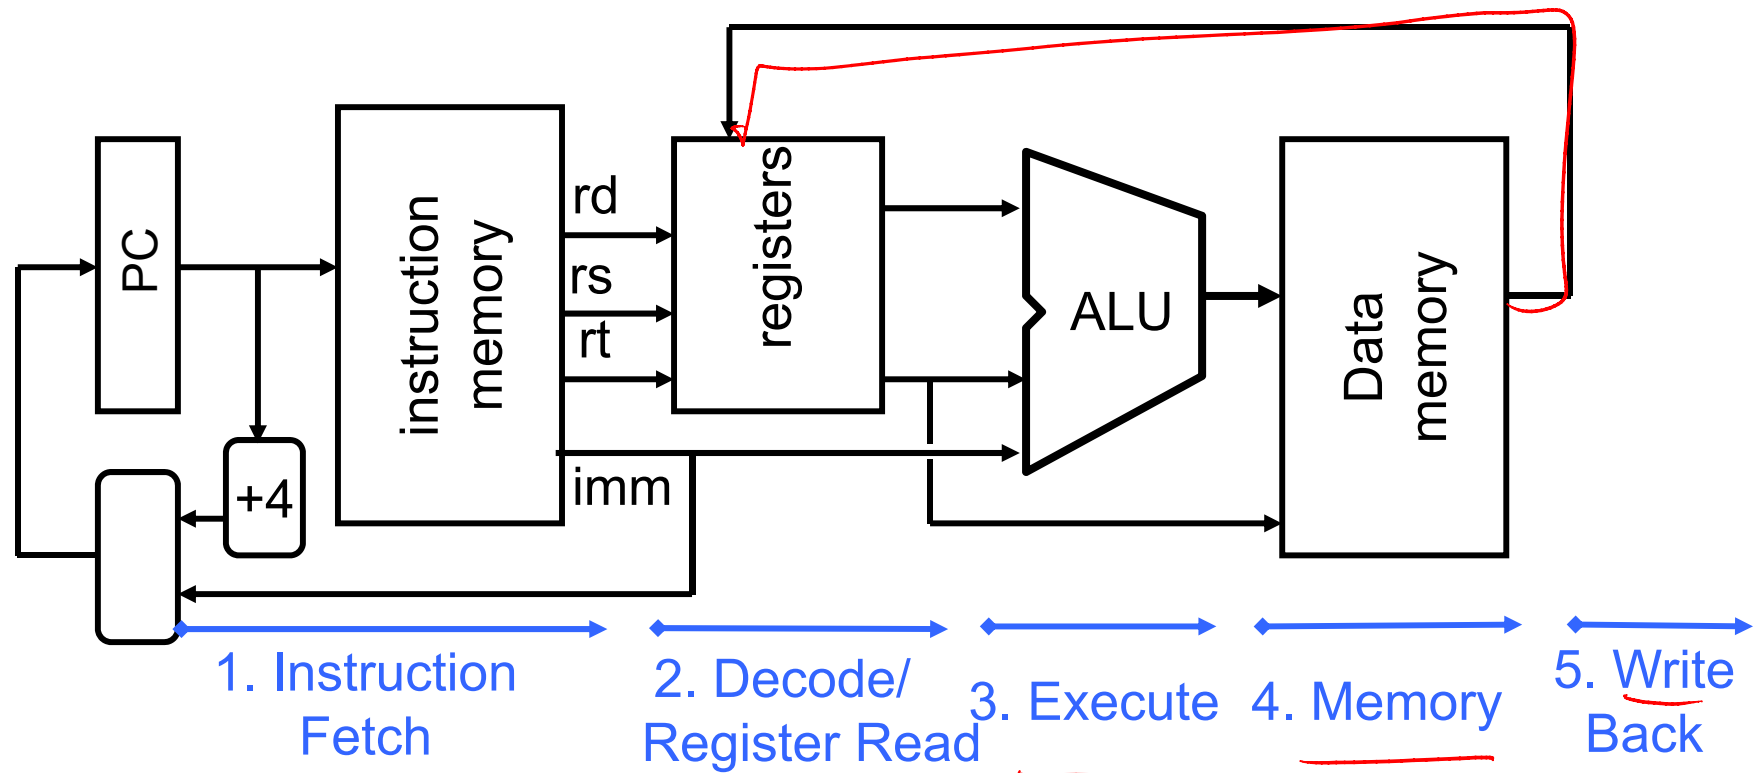
- Data Memory $\{R[rs] + \text{SignExt}[imm16]\} = R[rt]$



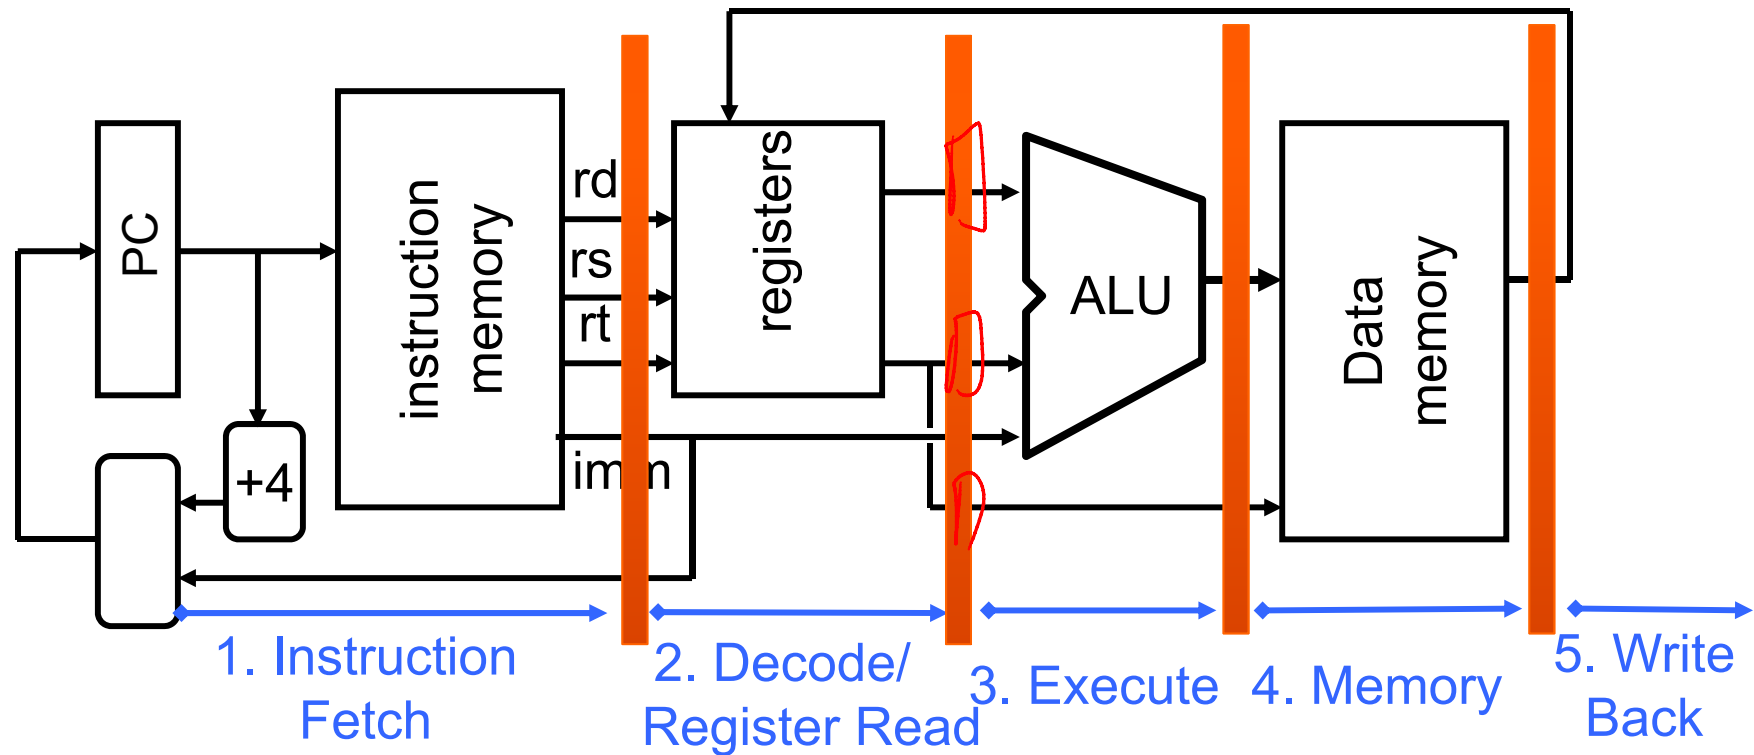
Steps in Executing MIPS

- 1) IFtch: Instruction Fetch, Increment PC
- 2) Dcd: Instruction Decode, Read Registers
- 3) Exec:
 - Mem-ref: Calculate Address
 - Arith-log: Perform Operation
- 4) Mem:
 - Load: Read Data from Memory
 - Store: Write Data to Memory
- 5) WB: Write Data Back to Register

Redrawn Single Cycle Datapath

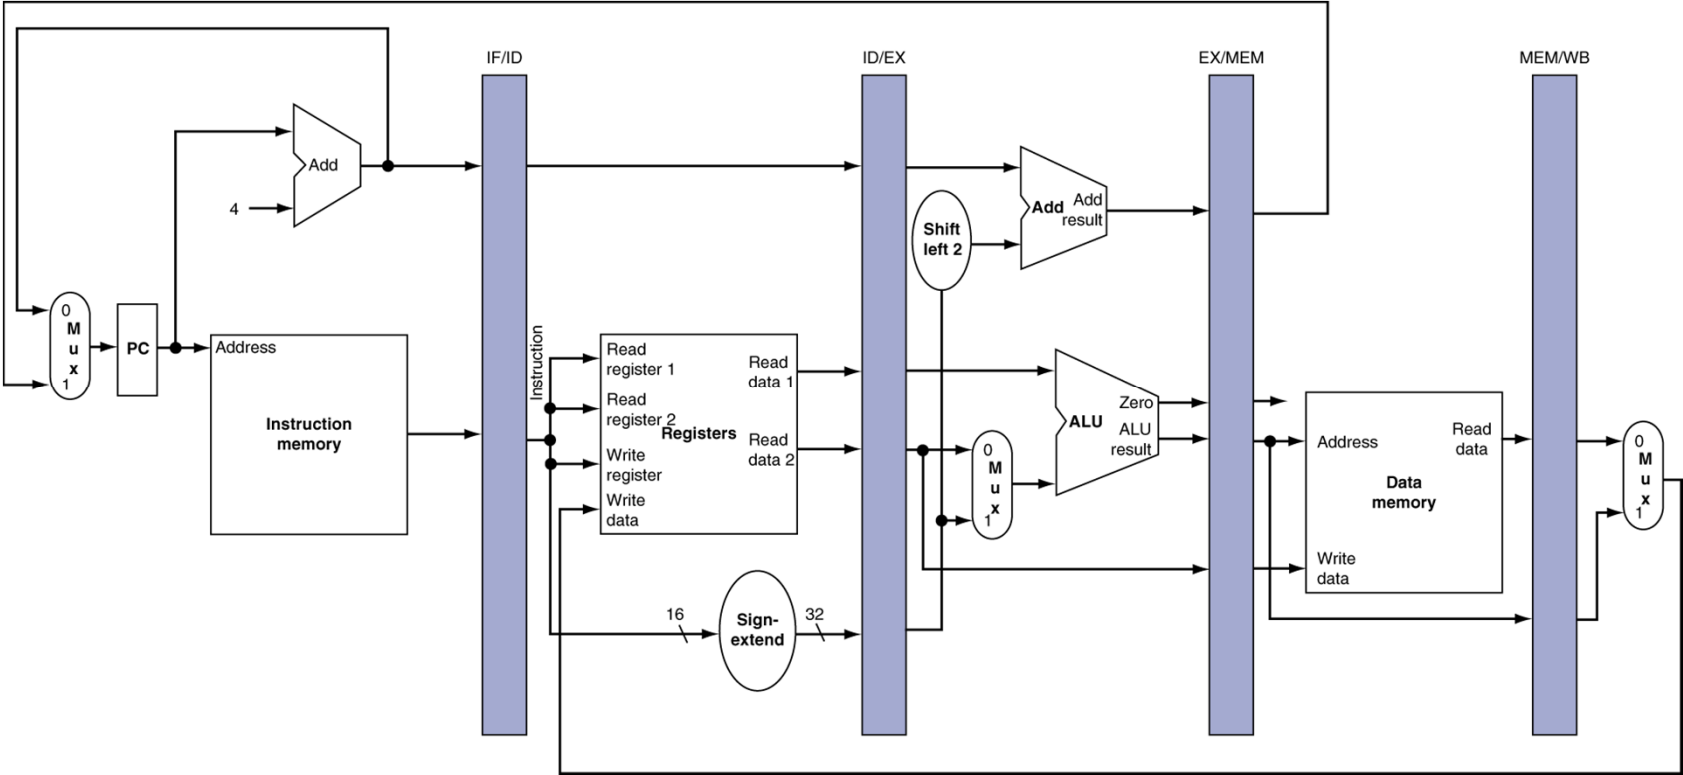


Pipeline registers

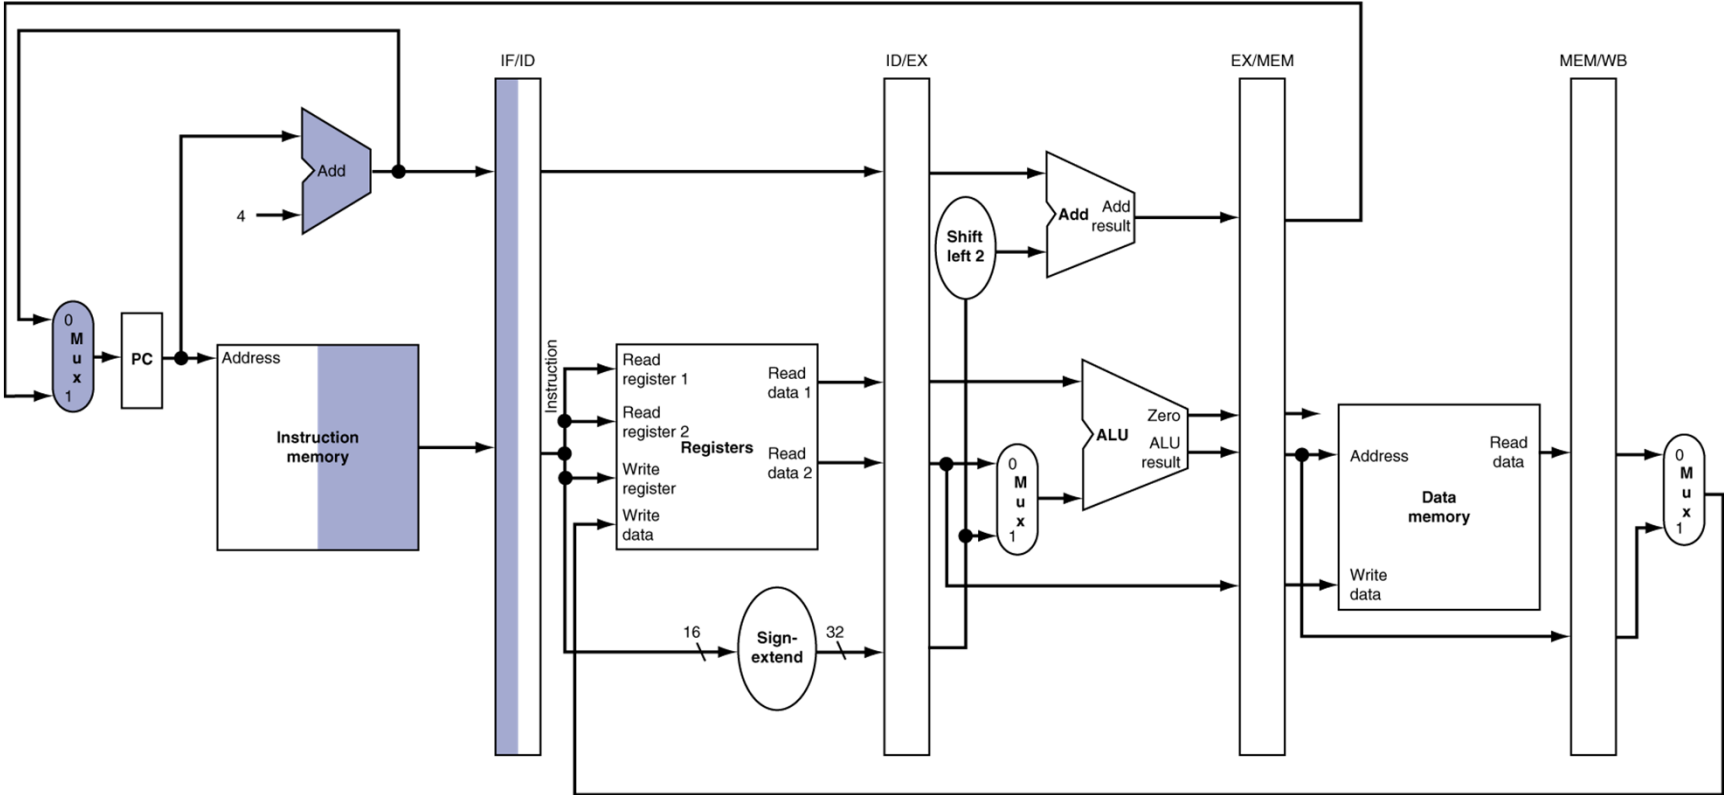



- Need registers between stages
 - To hold information produced in previous cycle

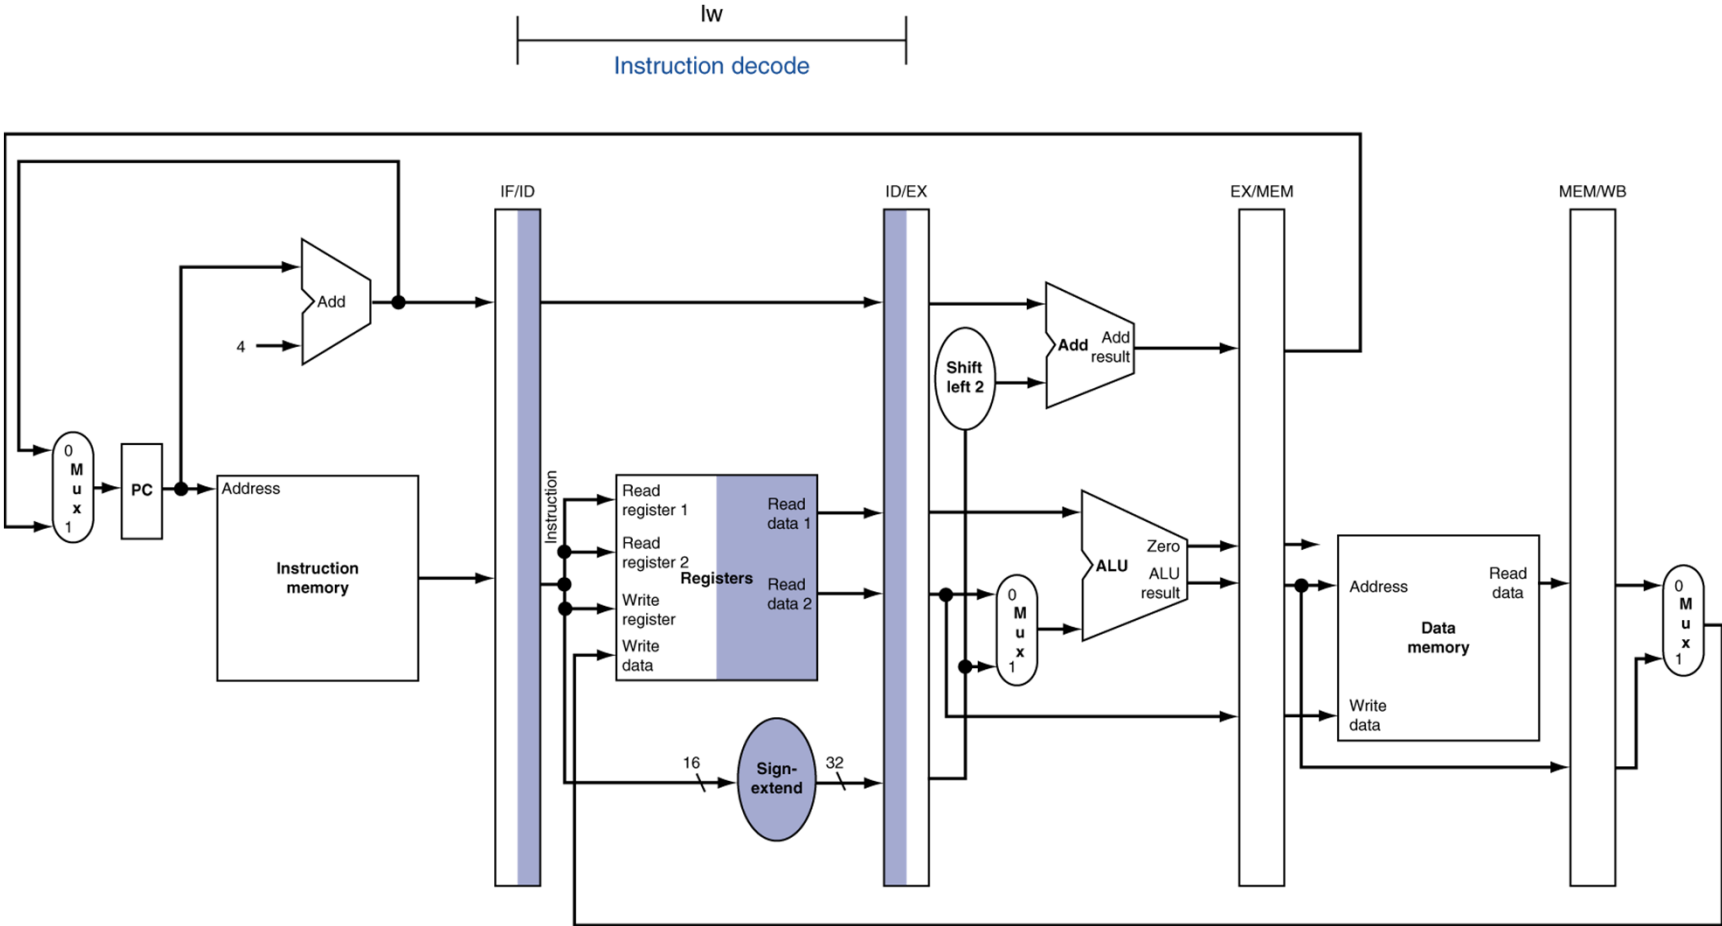
More Detailed Pipeline




IF for Load, Store, ...




ID for Load, Store, ...



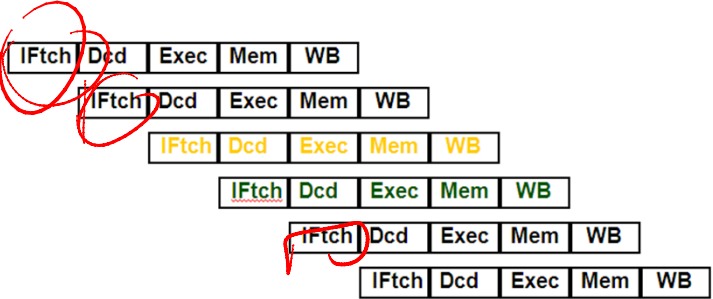
EX for Load



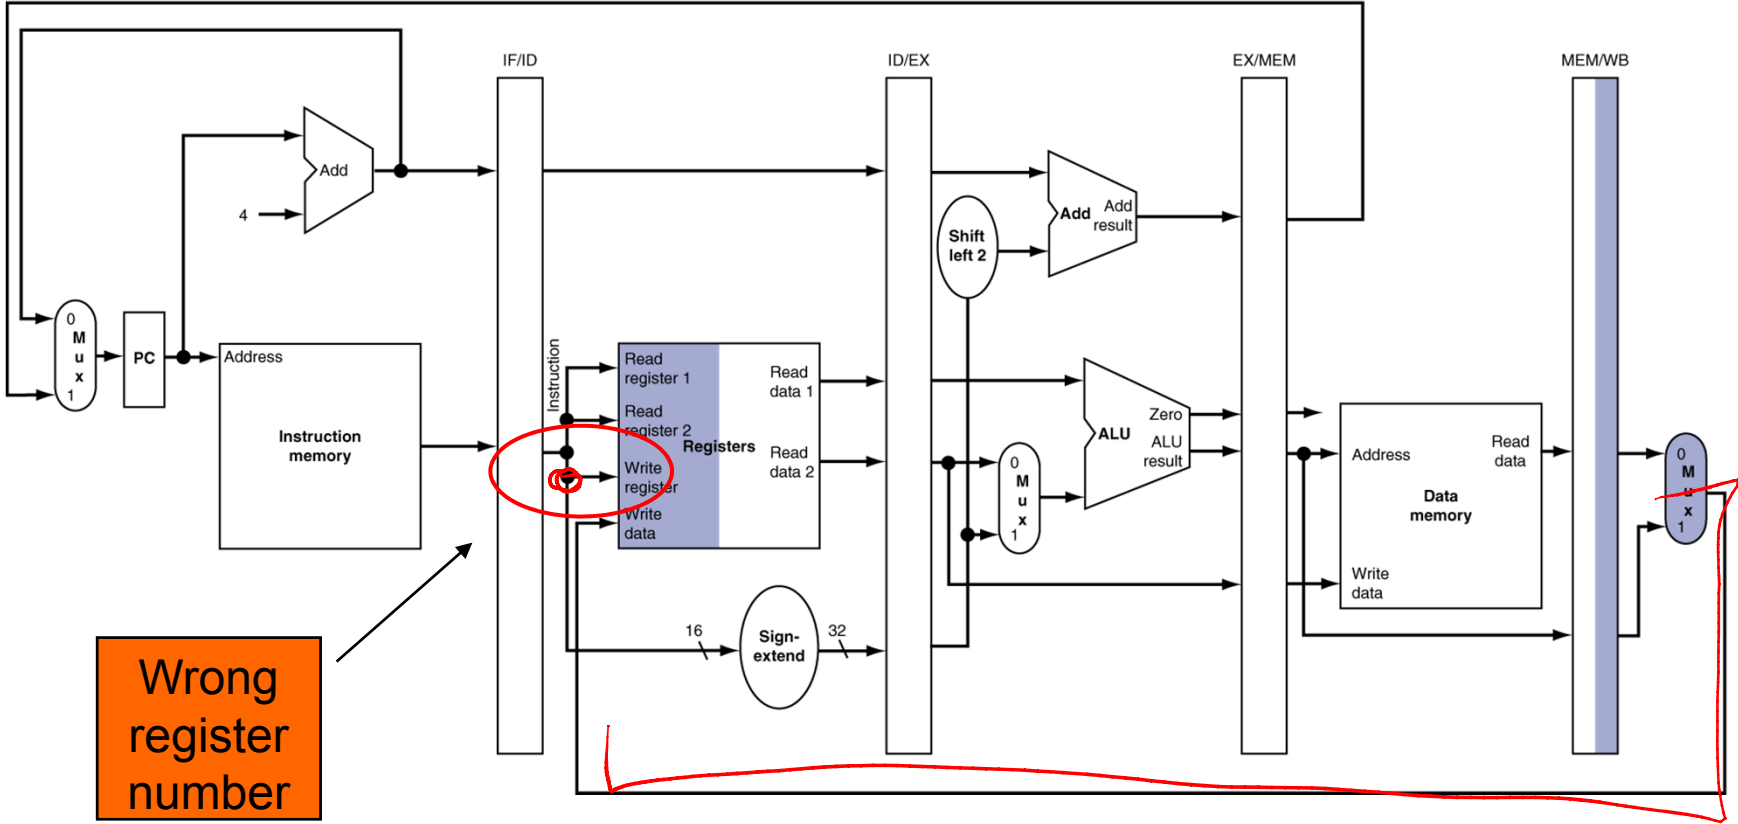
MEM for Load



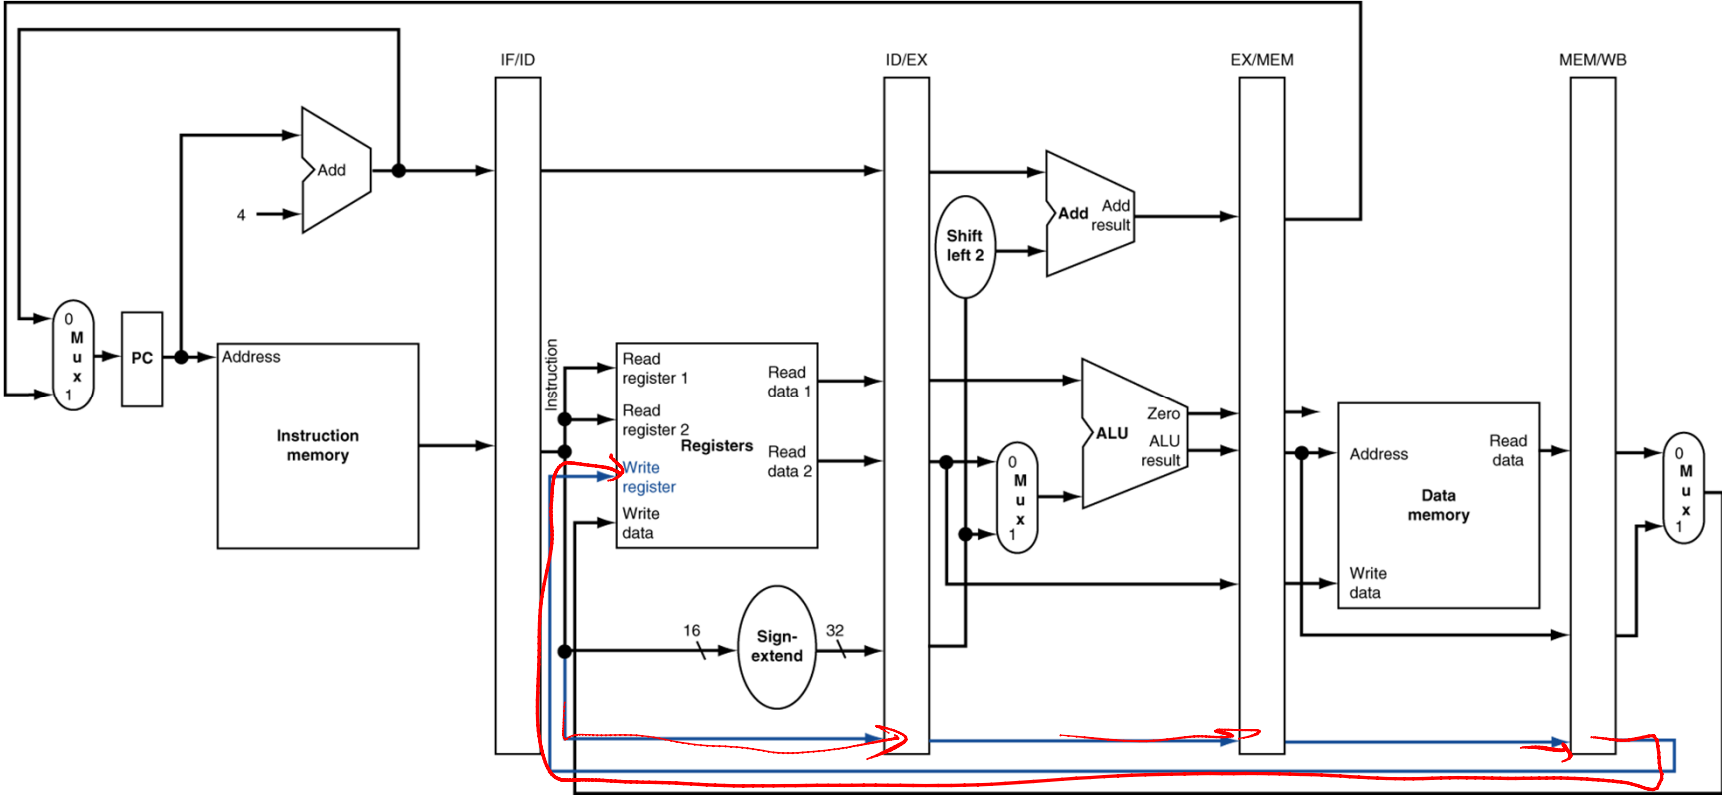
WB for Load



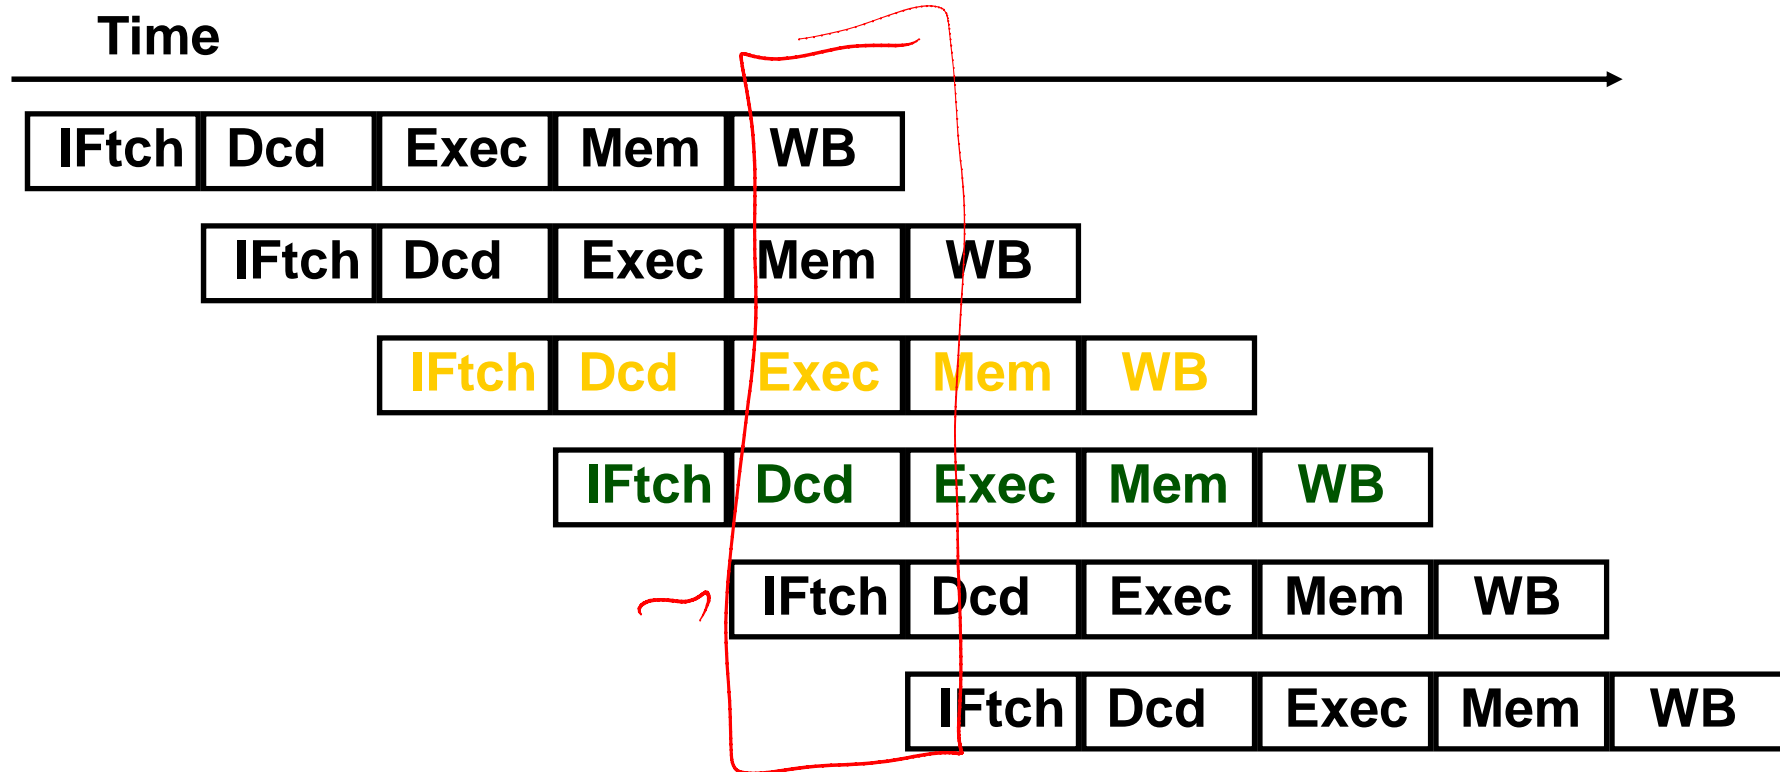
lw
Write back



Corrected Datapath for Load



Pipelined Execution Representation



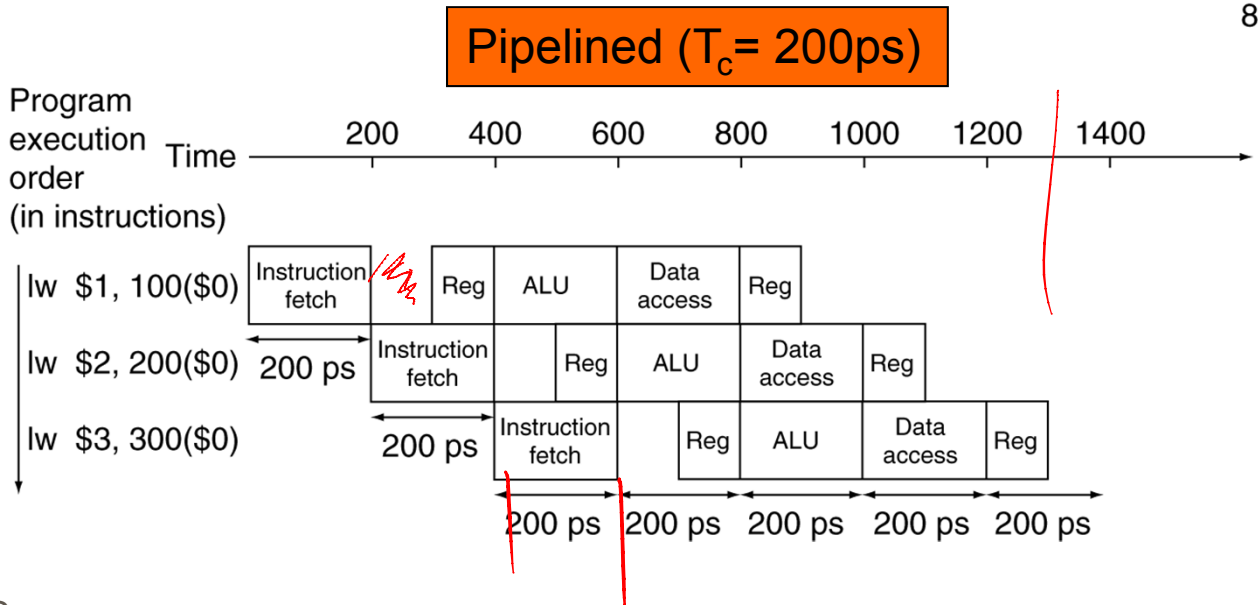
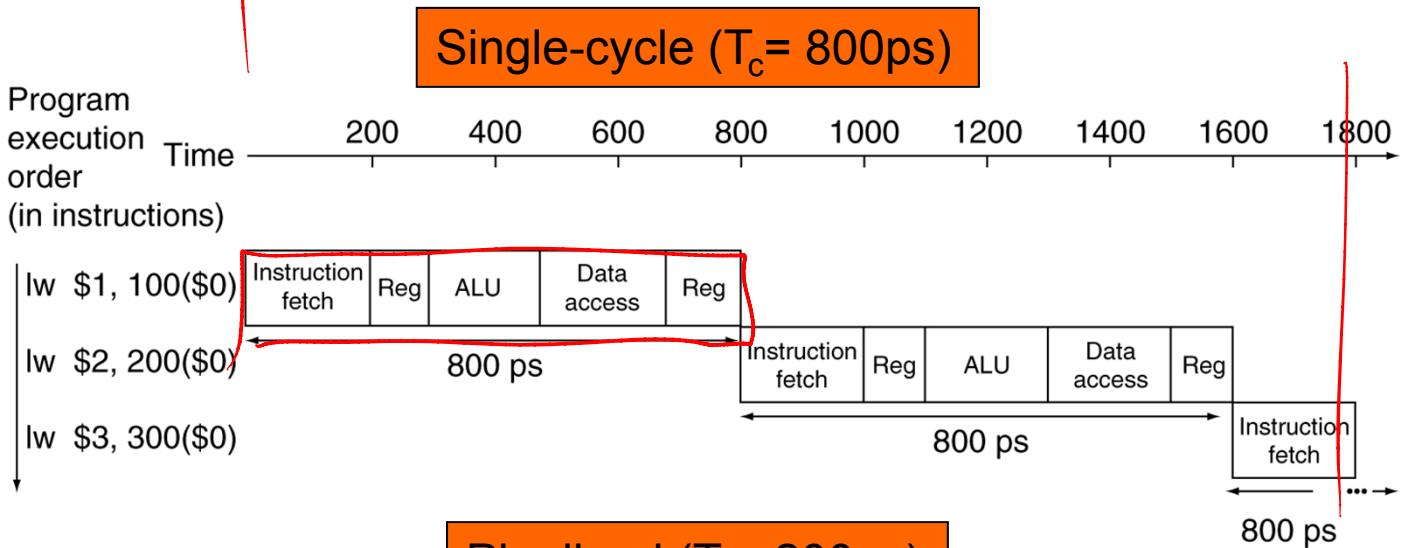
- Every instruction must take same number of steps, also called pipeline “stages”, so some will go idle sometimes

Pipeline Performance

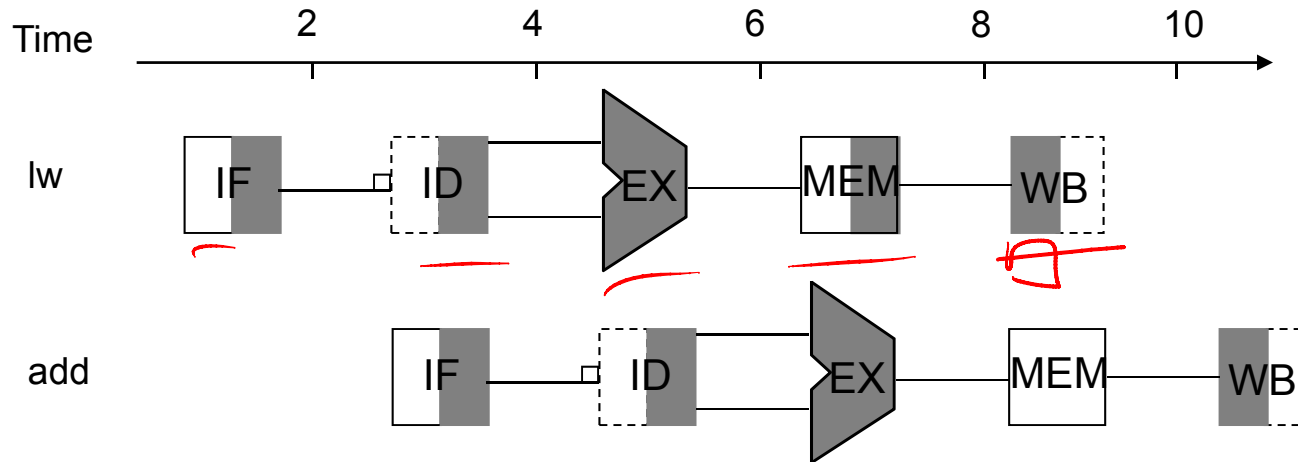
- Assume time for stages is
 - 100ps for register read or write
 - 200ps for other stages
- What is pipelined clock rate?
 - Compare pipelined datapath with single-cycle datapath

Instr	Instr fetch	Register read	ALU op	Memory access	Register write	Total time
lw	200ps	100 ps	200ps	200ps	100 ps	800ps
sw	200ps	100 ps	200ps	200ps		700ps
R-format	200ps	100 ps	200ps		100 ps	600ps
beq	200ps	100 ps	200ps			500ps

Pipeline Performance



Graphically Representing Pipelines




- Shading indicates the unit is being used by the instruction
- Shading on the **right half** of the register file (ID or WB) or memory means the element is being **read** in that stage
- Shading on the **left half** means the element is being **written** in that stage

Hazards

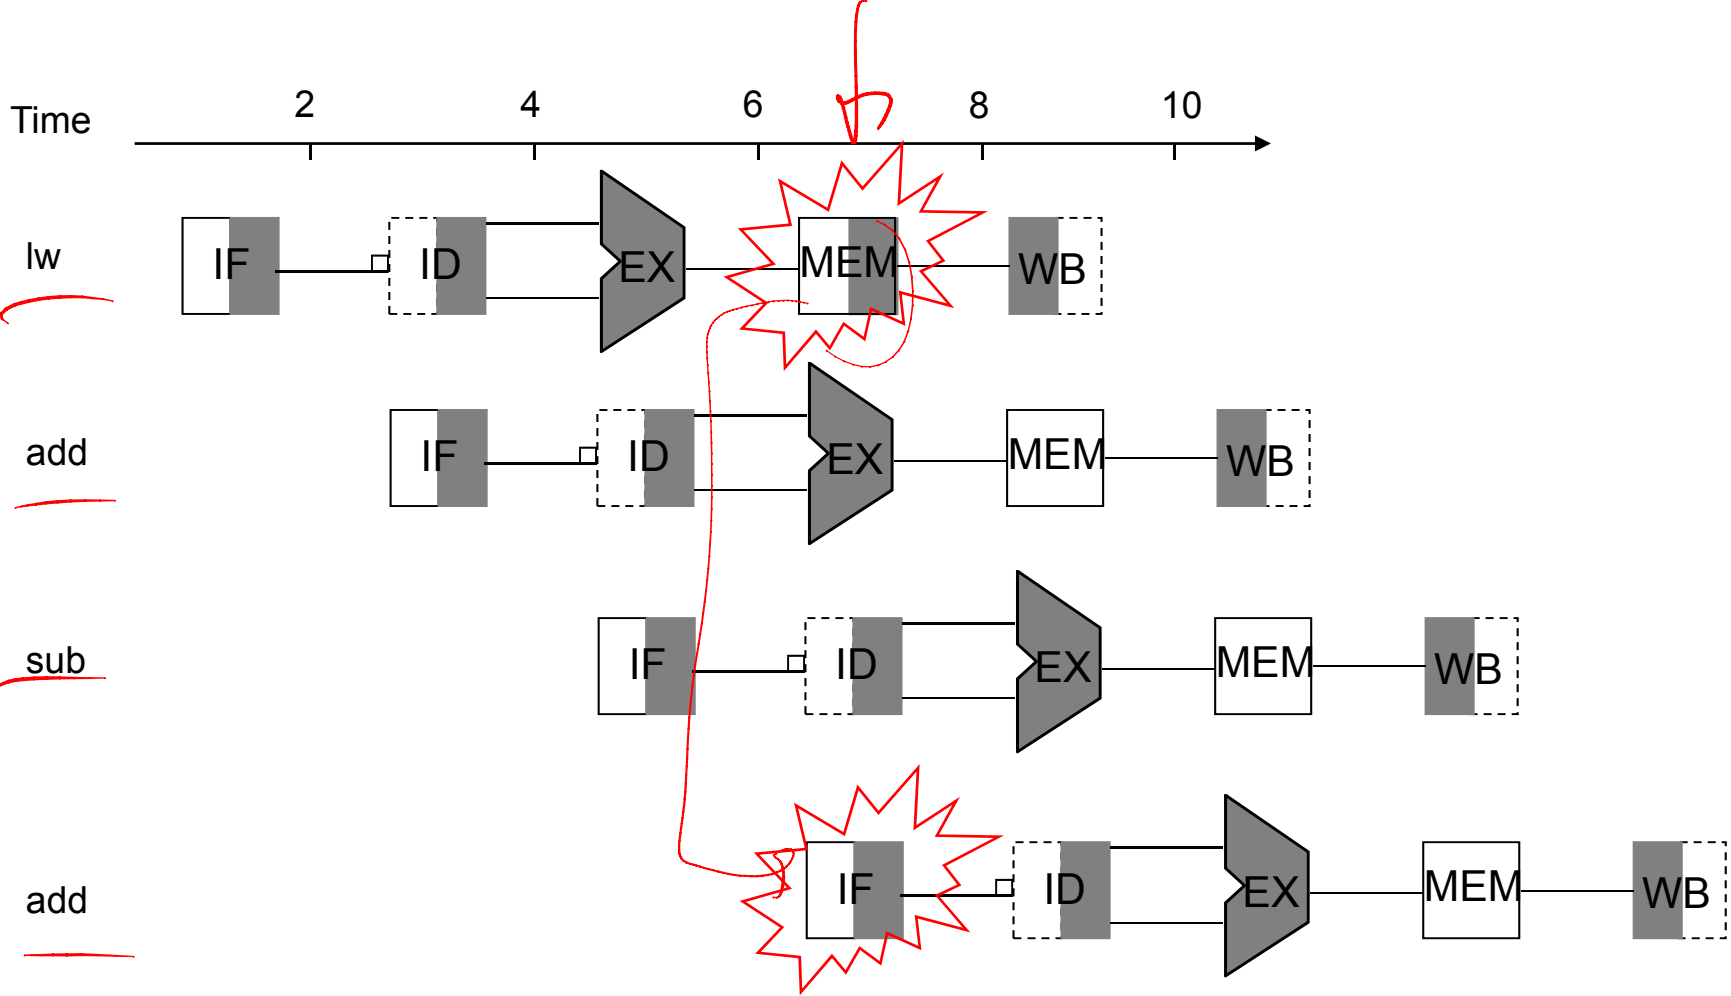
- It would be happy if we split the datapath into stages and the CPU works just fine
 - But, things are not that simple as you may expect
 - There are **hazards!**
- Situations that prevent starting the next instruction in the next cycle
 - **Structure hazards**
 - Conflict over the use of a resource at the same time
 - **Data hazard**
 - Data is not ready for the subsequent dependent instruction
 - **Control hazard**
 - Fetching the next instruction depends on the previous branch outcome

Structure Hazards

- Conflict over the use of a resource at the same time
- Suppose the MIPS CPU with a single memory
 - Load/store requires data access in MEM stage
 - Instruction fetch requires instruction access from the same memory
 - Instruction fetch would have to **stall** for that cycle
 - Would cause a **pipeline "bubble"**
- Hence, pipelined datapaths require separate instruction and data memories
 - Or separate instruction and data caches



Structure Hazards (Cont.)



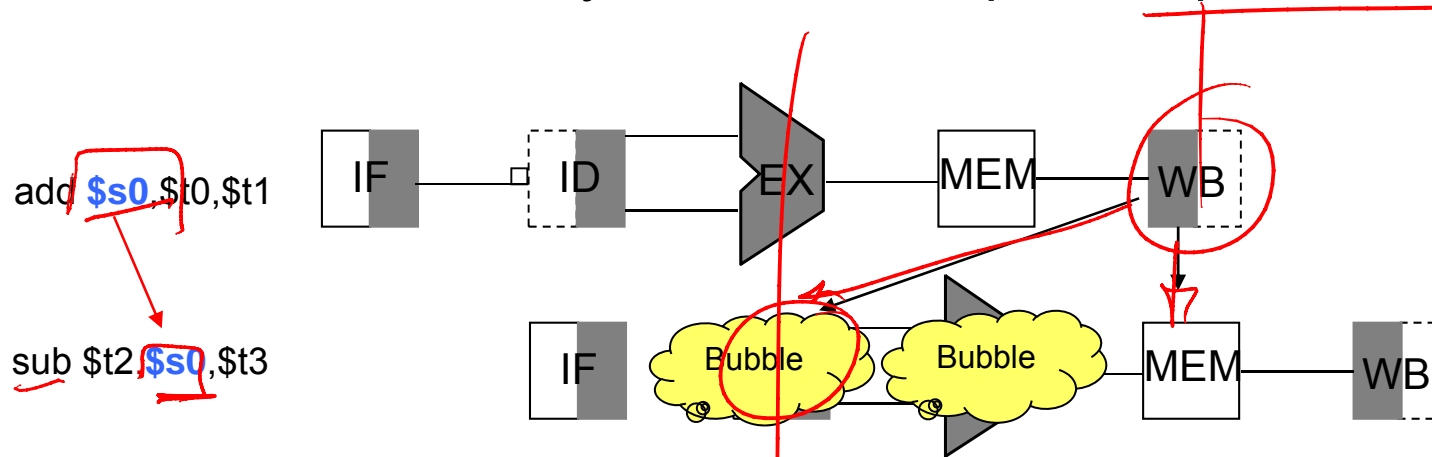
Need to separate instruction and data memory

Structural Hazard – reg read/write

- Two different solutions have been used:
 - 1) RegFile access is VERY fast: takes less than half the time of ALU stage
 - Write to Registers during first half of each clock cycle
 - Read from Registers during second half of each clock cycle
 - 2) Build RegFile with independent read and write ports
- **Result: can perform Read and Write during same clock cycle**


Data Hazards

- Data is not ready for the subsequent dependent instruction



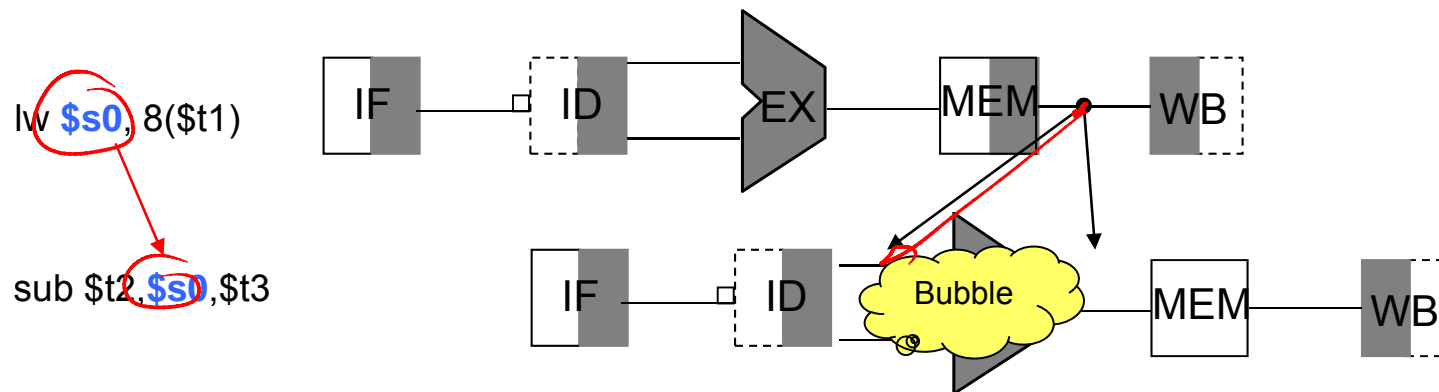
- To solve the data hazard problem, the pipeline needs to be stalled (typically referred to as “bubble”)
 - Then, performance is penalized
- A better solution?
 - **Forwarding** (or Bypassing)

Reducing Data Hazard - Forwarding



Data Hazard – Load-Use Case


- Can't always avoid stalls by forwarding
 - Can't forward backward in time!



- This bubble can be hidden by proper instruction scheduling

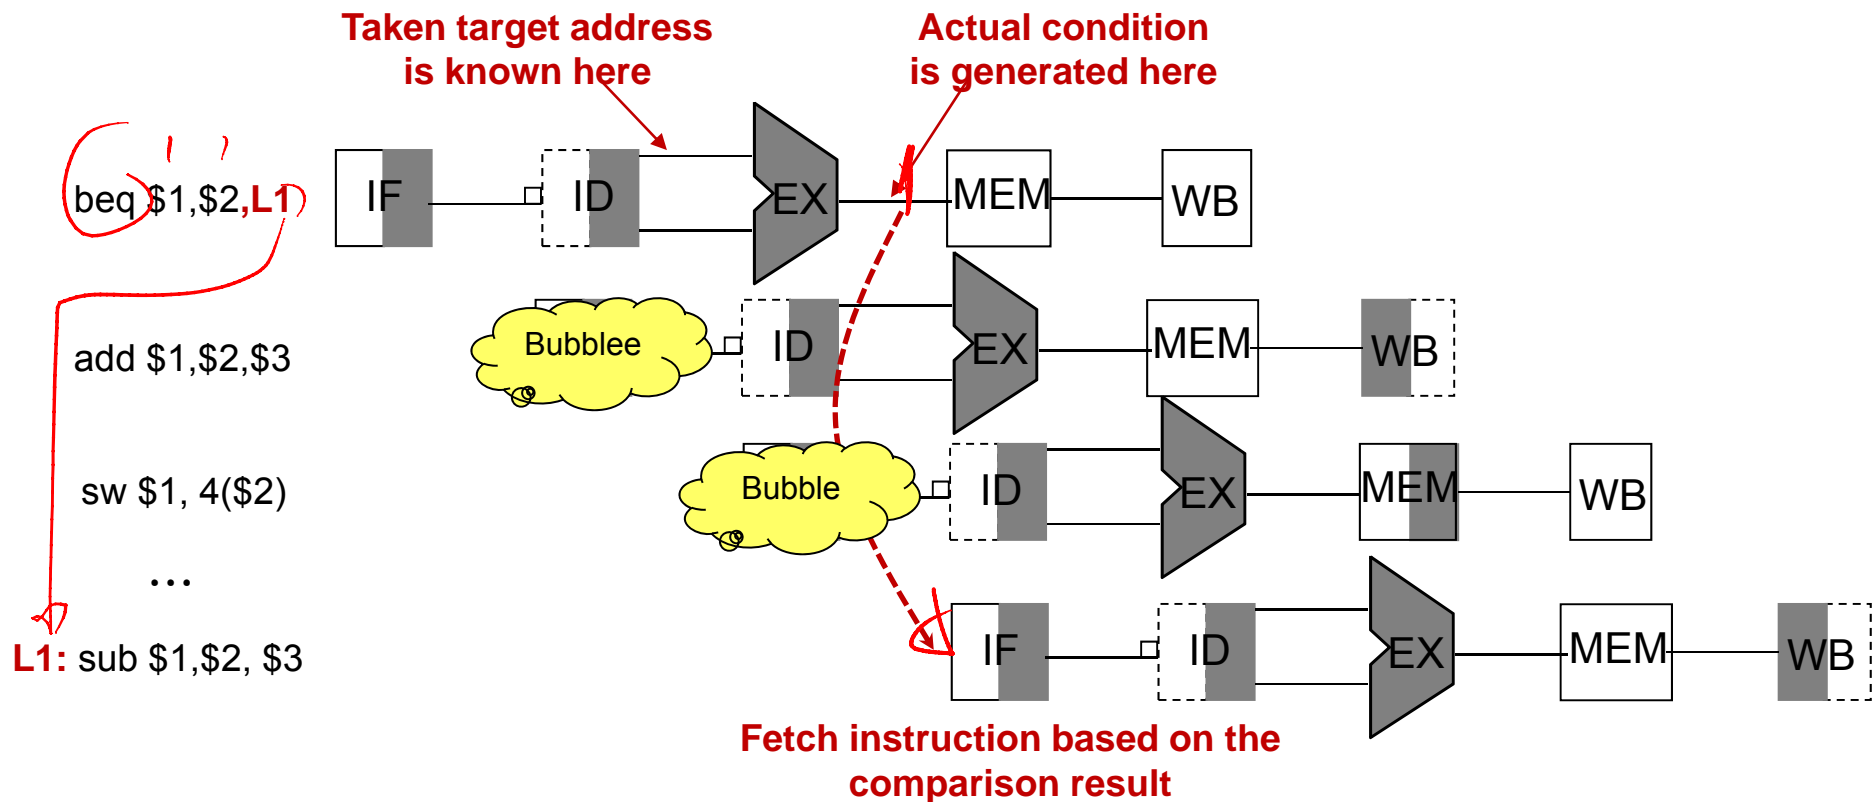
Code Scheduling to Avoid Stalls

- Reorder code to avoid use of load result in the next instruction
- C code for $A = B + E$; $C = B + F$;



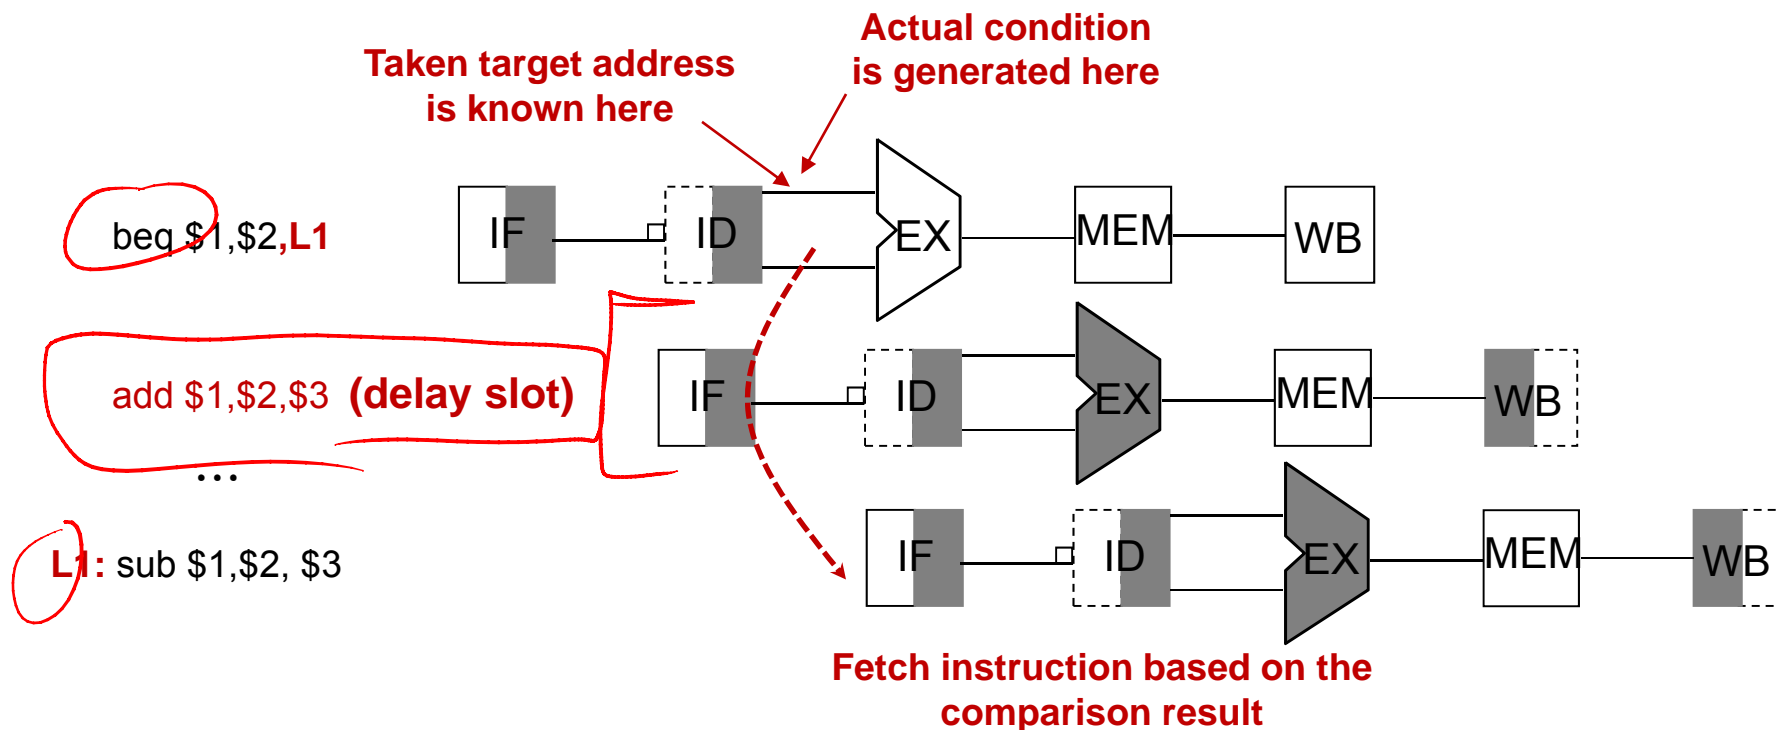
Control Hazard

- Branch determines the flow of instructions
- Fetching next instruction depends on branch outcome
 - Pipeline can't always fetch correct instruction
 - Branch instruction is still working on ID stage when fetching the next instruction




Delay Slot

- Branch instructions entail a “**delay slot**”
 - **Delayed branch** always executes the next sequential instruction, with the branch taking place after that one instruction delay
 - **Delay slot** is the slot right after a delayed branch instruction



Delay Slot (Cont.)

- Compiler needs to schedule a useful instruction in the delay slot, or fills it up with `nop` (no operation)



Can we do better?

```
bne $t0, $zero, L1
add $s1, $s2, $s3 // delay slot
addi $t1, $t1, 1
L1: addi $t1, $t1, 2
```

Fill the delay slot with
a useful and valid
instruction

Pipeline Summary

- Pipelining improves performance by increasing instruction **throughput**
 - Executes multiple instructions in parallel
- Pipelining is subject to hazards
 - Structure, data, control hazards
- Instruction set design affects the complexity of the pipeline implementation

Embedded Processors : examples

CISC	RISC
68000 series	Sparc
X86 family	AMD 29000
PDP-11	MIPS
VAX	SuperH
IBM 370	<u>PowerPC</u>
	Arm