

The diagram on the left shows three circles. The bottom circle contains a transistor circuit diagram. An arrow points from it to the middle circle, which contains a purple square labeled 'CPU'. Another arrow points from the middle circle to the top circle, which contains a flowchart with a diamond, a rectangle, and a green bar.

## Chapter 5 The LC-3

Original slides from Gregory Byrd, North Carolina State University  
Modified slides by C. Wilcox, S. Rajopadhye Colorado State University

Copyright © The McGraw-Hill Companies, Inc. Permission required for reproduction or display.

## Computing Layers

The diagram on the left is identical to the one in the first slide, showing the progression from hardware (transistor) to CPU to software (flowchart).

Problems

-----

Algorithms

-----

Language

-----

Instruction Set Architecture ←

-----

Microarchitecture

-----

Circuits

-----

Devices

CS270 - Fall 2013 - Colorado State University 2

## Instruction Set Architecture

- ISA = All of the **programmer-visible** components and operations of the computer
  - **memory organization**
    - address space -- how many locations can be addressed?
    - addressability -- how many bits per location?
  - **register set**
    - how many? what size? how are they used?
  - **instruction set**
    - opcodes
    - data types
    - addressing modes
- ISA provides all information needed for someone that wants to write a program in **machine language**
  - or translate from a high-level language to machine language.

## LC-3 Overview: Memory and Registers

- **Memory**
  - address space: **2<sup>16</sup>** locations (16-bit addresses)
  - addressability: **16 bits**
- **Registers**
  - temporary storage, accessed in a single machine cycle
    - accessing memory takes longer than a single cycle
  - eight general-purpose registers: **R0 - R7**
    - each **16 bits wide**
    - how many bits to uniquely identify a register?
  - other registers
    - not directly addressable, but used by (and affected by) instructions
    - **PC** (program counter), **condition codes**

## LC-3 Overview: Instruction Set

- **Opcodes**
  - 15 opcodes, 3 types of instructions
  - **Operate**: ADD, AND, NOT
  - **Data movement**: LD, LDI, LDR, LEA, ST, STR, STI
  - **Control**: BR, JSR/JSRR, JMP, RTI, TRAP
  - some opcodes set/clear *condition codes*, based on result:
    - N = negative, Z = zero, P = positive (> 0)
- **Data Types**
  - 16-bit 2's complement integer
- **Addressing Modes**
  - How is the location of an operand specified?
  - non-memory addresses: *immediate, register*
  - memory addresses: *PC-relative, indirect, base+offset*


## Operate Instructions

- Only three operations: **ADD, AND, NOT**
- Source and destination operands are **registers**
  - These instructions *do not* reference memory.
  - ADD and AND can use “immediate” mode, where one operand is hard-wired into the instruction.
- Will show *dataflow diagram* with each instruction.
  - illustrates *when* and *where* data moves to accomplish the desired operation


# NOT (Register)



Note: Src and Dst could be the same register.



# ADD/AND (Register)

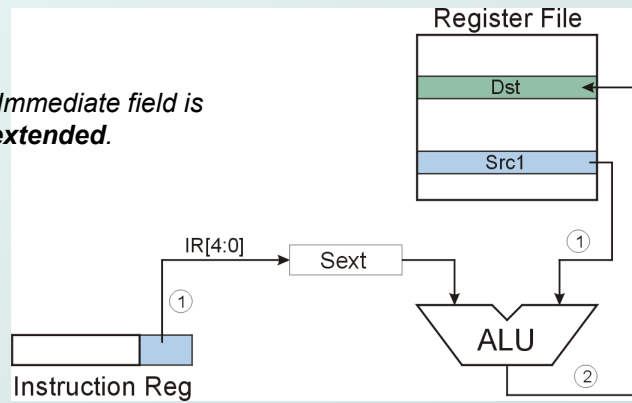


## ADD/AND (Immediate)

this one means "immediate mode"



Note: Immediate field is **sign-extended**.



## Using Operate Instructions

- With only ADD, AND, NOT...
  - How do we subtract?
  - How do we OR?
  - How do we copy from one register to another?
  - How do we initialize a register to zero?


## Data Movement Instructions

- Load -- read data **from memory to register**
  - **LD**: PC-relative mode
  - **LDR**: base+offset mode
  - **LDI**: indirect mode
- Store -- write data **from register to memory**
  - **ST**: PC-relative mode
  - **STR**: base+offset mode
  - **STI**: indirect mode
- Load effective address -- compute address, save in register
  - **LEA**: immediate mode
  - *does not access memory*

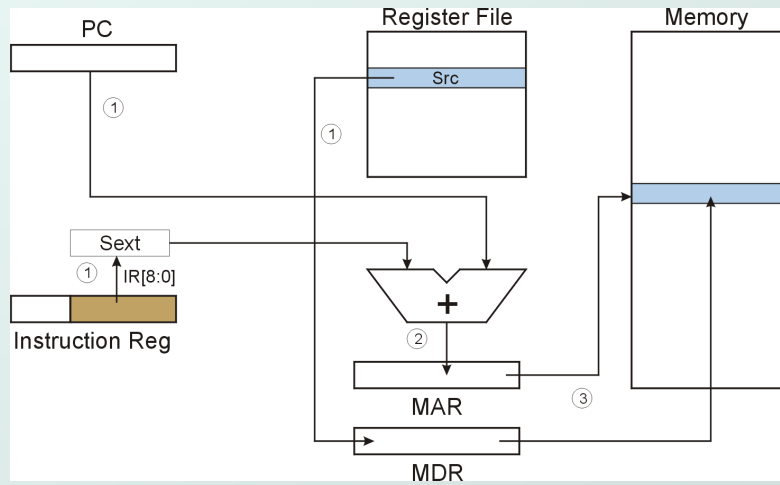
## PC-Relative Addressing Mode

- Want to specify address directly in the instruction
  - But an address is 16 bits, and so is an instruction!
  - After subtracting 4 bits for opcode and 3 bits for register, we have 9 bits available for address.
- **Solution:**
  - Use the 9 bits as a **signed offset** from the current PC.
- 9 bits:  $-256 \leq \text{offset} \leq +255$
- Can form address such that:  $\text{PC} - 256 \leq X \leq \text{PC} + 255$ 
  - Remember that PC is incremented as part of the FETCH phase;
  - This is done before the EVALUATE ADDRESS stage.

## LD (PC-Relative)



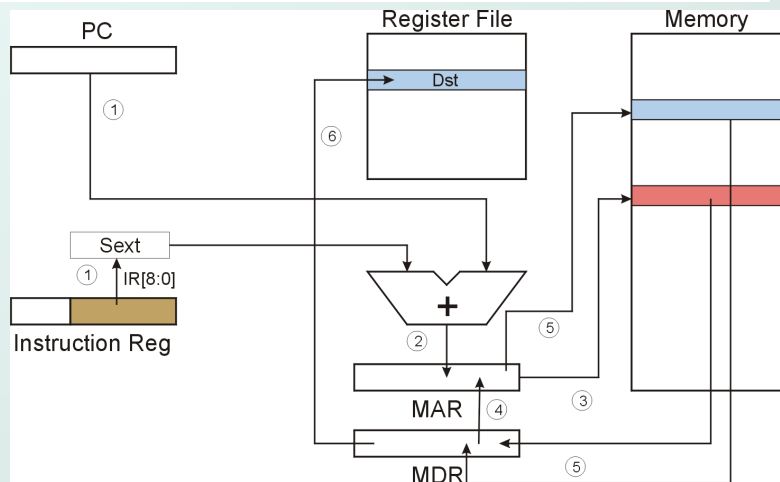
## ST (PC-Relative)



## Indirect Addressing Mode

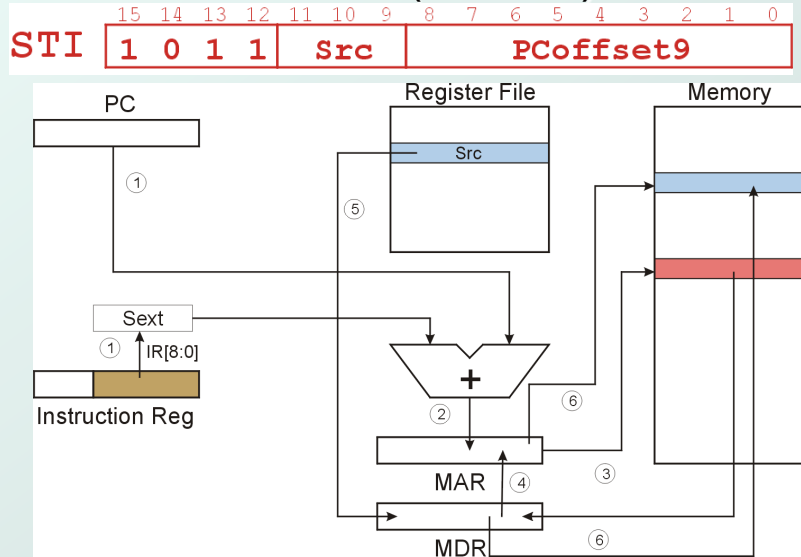
- With PC-relative mode, can only address data within 256 words of the instruction.
  - What about the rest of memory?
- **Solution #1:**
  - **Read address from memory location, then load/store to that address.**
- First address is generated from PC and IR (just like PC-relative addressing), then content of that address is used as target for load/store.

## LDI (Indirect)





## STI (Indirect)



CS270 - Fall 2013 - Colorado State University

17


## Base + Offset Addressing Mode

- With PC-relative mode, can only address data within 256 words of the instruction.
  - What about the rest of memory?
- **Solution #2:**
  - **Use a register to generate a full 16-bit address.**
- 4 bits for opcode, 3 for src/dest register, 3 bits for **base** register -- remaining 6 bits are used as a **signed offset**.
  - Offset is **sign-extended** before adding to base register.

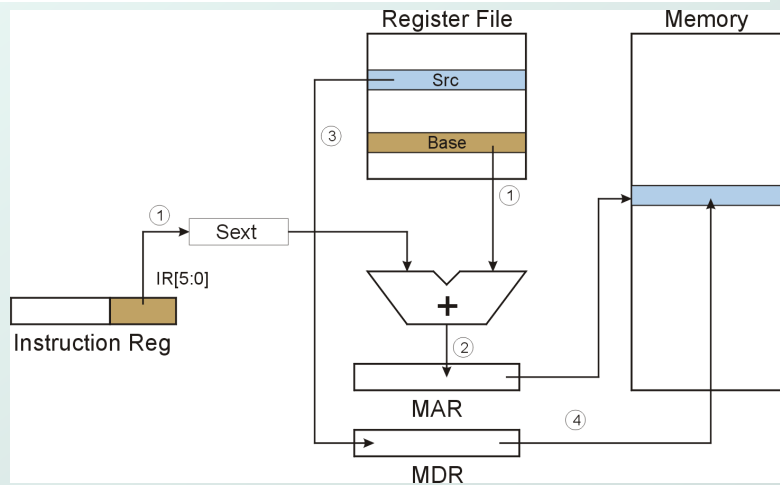
CS270 - Fall 2013 - Colorado State University

18

## LDR (Base+Offset)



## STR (Base+Offset)

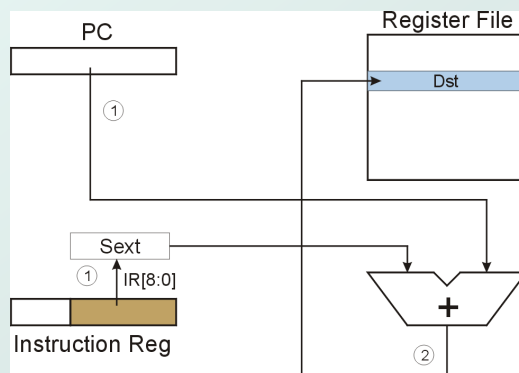


## Load Effective Address

- Computes address like PC-relative (PC plus signed offset) and **stores the result into a register.**

**Note:** The address is stored in the register, not the contents of the memory location.

## LEA (Immediate)



## Example

Address	Instruction	Comments
x30F6	1 1 1 0 0 0 1 1 1 1 1 1 1 1 0 1	$R1 \leftarrow PC - 3 = x30F4$
x30F7	0 0 0 1 0 1 0 0 0 1 1 0 1 1 1 0	$R2 \leftarrow R1 + 14 = x3102$
x30F8	0 0 1 1 0 1 0 1 1 1 1 1 1 0 1 1	$M[PC - 5] \leftarrow R2$ $M[x30F4] \leftarrow x3102$
x30F9	0 1 0 1 0 1 0 0 1 0 1 0 0 0 0 0	$R2 \leftarrow 0$
x30FA	0 0 0 1 0 1 0 0 1 0 1 0 0 1 0 1	$R2 \leftarrow R2 + 5 = 5$
x30FB	0 1 1 1 0 1 0 0 0 1 0 0 1 1 1 0	$M[R1+14] \leftarrow R2$ $M[x3102] \leftarrow 5$
x30FC	1 0 1 0 0 1 1 1 1 1 1 1 0 1 1 1 <i>opcode</i>	$R3 \leftarrow M[M[x30F4]]$ $R3 \leftarrow M[x3102]$ $R3 \leftarrow 5$

## Register Transfer Notation/Level

- Used to describe the operational behavior of digital circuits
  - Cycle by cycle or at a more “macro” level
    - $R1 \leftarrow PC - 3 = x30F4$
    - $DstReg \leftarrow Value$ , using other regs or memory and ops
- Also names which control signals are on (i.e., 1) during a cycle. By default signals not named are off
- Control signals are the critical elements, everything else can be inferred.
  - More of this in recitation

## Control Instructions

- Used to alter the sequence of instructions (by changing the Program Counter)
- **Conditional Branch**
  - branch is *taken* if a specified condition is true
    - signed offset is added to PC to yield new PC
  - else, the branch is *not taken*
    - PC is not changed, points to the next instruction
- **Unconditional Branch (or Jump)**
  - always changes the PC
- **TRAP**
  - changes PC to the address of an OS “service routine”
  - routine will return control to the next instruction (after the TRAP)

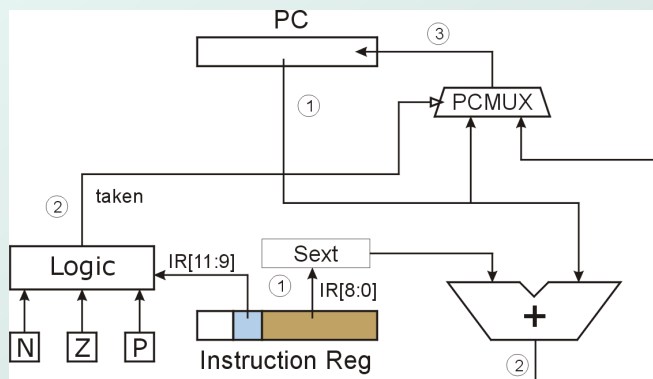
## Condition Codes

- LC-3 has three **condition code** registers:
  - N** -- negative
  - Z** -- zero
  - P** -- positive (greater than zero)
- Set by any instruction that writes a value to a register  
(ADD, AND, NOT, LD, LDR, LDI, LEA)
- Exactly one will be set at all times
  - Based on the last instruction that altered a register

## Branch Instruction

- Branch specifies one or more condition codes.
- If the set bit is specified, the branch is taken.
  - PC-relative addressing:  
**target address** is made by adding signed offset (IR[8:0]) to current PC.
  - Note: PC has already been incremented by FETCH stage.
  - Note: Target must be within 256 words of BR instruction.
- If the branch is not taken, the next sequential instruction is executed.

## BR (PC-Relative)

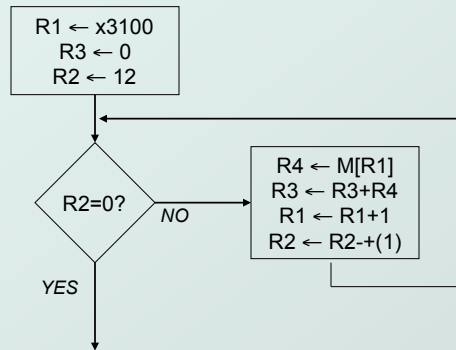


What happens if bits [11:9] are all zero? All one?

## Using Branch Instructions

- Compute sum of 12 integers.

Numbers start at location x3100. Program starts at location x3000.




## Sample Program

Address	Instruction	Comments
x3000	1 1 1 0 0 0 1 0 1 1 1 1 1 1 1 1	<b>R1 ← x3100 (PC+0xFF)</b>
x3001	0 1 0 1 0 1 1 0 1 1 1 0 0 0 0 0	<b>R3 ← 0</b>
x3002	0 1 0 1 0 1 0 0 1 0 1 0 0 0 0 0	<b>R2 ← 0</b>
x3003	0 0 0 1 0 1 0 0 1 0 1 0 1 1 0 0	<b>R2 ← 12</b>
x3004	0 0 0 0 0 1 0 0 0 0 0 0 0 1 0 1	<b>If Z, goto x300A (PC+5)</b>
x3005	0 1 1 0 1 0 0 0 0 1 0 0 0 0 0 0	<b>Load next value to R4</b>
x3006	0 0 0 1 0 1 1 0 1 1 0 0 0 0 0 1	<b>Add to R3</b>
x3007	0 0 0 1 0 0 1 0 0 1 1 0 0 0 0 1	<b>Increment R1 (pointer)</b>
X3008	0 0 0 1 0 1 0 0 1 0 1 1 1 1 1 1	<b>Decrement R2 (counter)</b>
x3009	0 0 0 0 1 1 1 1 1 1 1 1 1 0 1 0	<b>Goto x3004 (PC-6)</b>

## JMP (Register)

- Jump is an unconditional branch -- always taken.
  - Target address is the contents of a register.
  - Allows any target address.

**JMP**    15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0  
 1 1 0 0 0 0 0 Base 0 0 0 0 0 0



## TRAP

**TRAP**    15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0  
 1 1 1 1 0 0 0 0 trapvect8

- Calls a **service routine**, identified by 8-bit “trap vector.”

vector	routine
x23	input a character from the keyboard
x21	output a character to the monitor
x25	halt the program

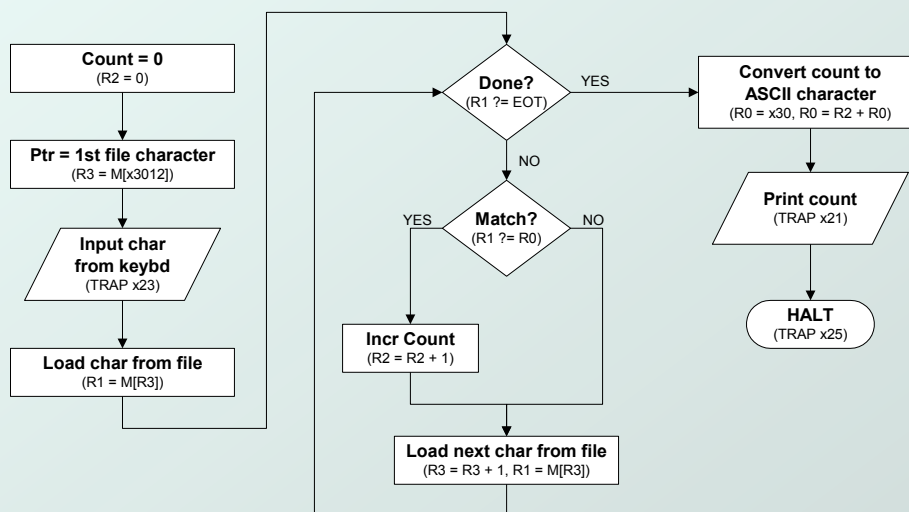
- When routine is done, PC is set to the instruction following TRAP.
  - We’ ll talk about how this works later.



## Another Example

- **Count the occurrences of a character in a file**
  - Program begins at location x3000
  - Read character from keyboard
  - Load each character from a “file”
    - File is a sequence of memory locations
    - Starting address of file is stored in the memory location immediately after the program
  - If file character equals input character, increment counter
  - End of file is indicated by an ASCII value: **EOT (x04)**
  - At the end, print the number of characters and halt (assume there will be less than 10 occurrences of the character)
- A special character used to indicate the end of a sequence is often called a **sentinel**.
  - Useful when you don't know ahead of time how many times to execute a loop.

## Flow Chart

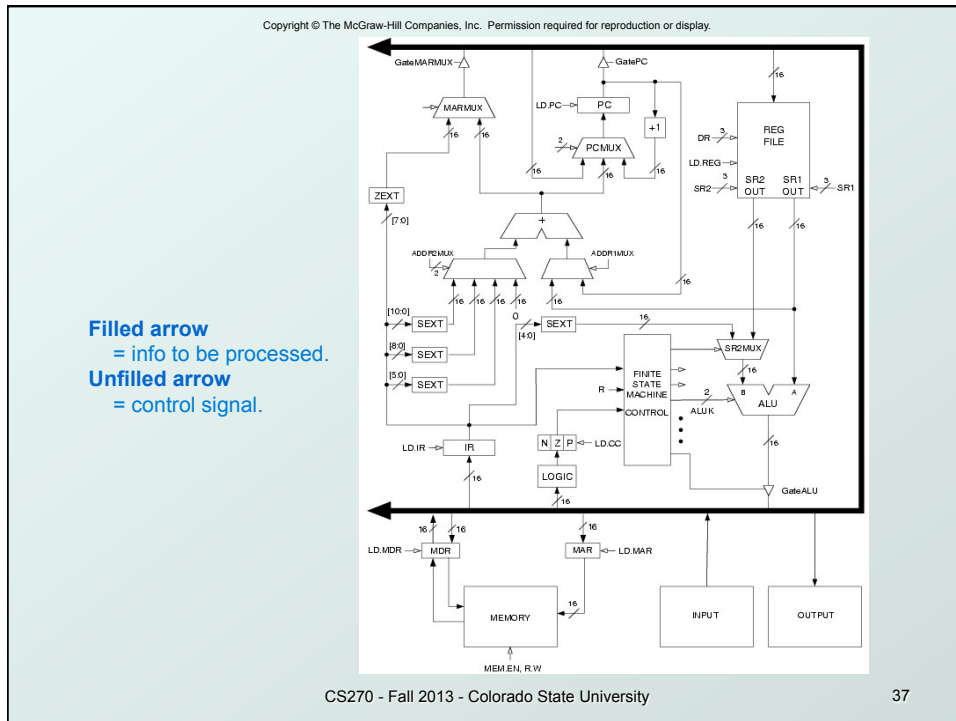


## Program (1 of 2)

Address	Instruction	Comments
x3000	0 1 0 1 0 1 0 0 1 0 1 0 0 0 0 0	<b><math>R2 \leftarrow 0</math> (counter)</b>
x3001	0 0 1 0 0 1 1 0 0 0 0 1 0 0 0 0	<b><math>R3 \leftarrow M[x3102]</math> (ptr)</b>
x3002	1 1 1 1 0 0 0 0 0 0 1 0 0 0 1 1	<b>Input to R0 (TRAP x23)</b>
x3003	0 1 1 0 0 0 1 0 1 1 0 0 0 0 0 0	<b><math>R1 \leftarrow M[R3]</math></b>
x3004	0 0 0 1 1 0 0 0 0 1 1 1 1 1 0 0	<b><math>R4 \leftarrow R1 - 4</math> (EOT)</b>
x3005	0 0 0 0 0 1 0 0 0 0 0 0 1 0 0 0	<b>If Z, goto x300E</b>
x3006	1 0 0 1 0 0 1 0 0 1 1 1 1 1 1 1	<b><math>R1 \leftarrow \text{NOT } R1</math></b>
x3007	0 0 0 1 0 0 1 0 0 1 1 0 0 0 0 1	<b><math>R1 \leftarrow R1 + 1</math></b>
X3008	0 0 0 1 0 0 1 0 0 1 0 0 0 0 0 0	<b><math>R1 \leftarrow R1 + R0</math></b>
x3009	0 0 0 0 1 0 1 0 0 0 0 0 0 0 0 1	<b>If N or P, goto x300B</b>

## Program (2 of 2)

Address	Instruction	Comments
x300A	0 0 0 1 0 1 0 0 1 0 1 0 0 0 0 1	<b><math>R2 \leftarrow R2 + 1</math></b>
x300B	0 0 0 1 0 1 1 0 1 1 1 0 0 0 0 1	<b><math>R3 \leftarrow R3 + 1</math></b>
x300C	0 1 1 0 0 0 1 0 1 1 0 0 0 0 0 0	<b><math>R1 \leftarrow M[R3]</math></b>
x300D	0 0 0 0 1 1 1 1 1 1 1 1 0 1 1 0	<b>Goto x3004</b>
x300E	0 0 1 0 0 0 0 0 0 0 0 0 0 1 0 0	<b><math>R0 \leftarrow M[x3013]</math></b>
x300F	0 0 0 1 0 0 0 0 0 0 0 0 0 0 1 0	<b><math>R0 \leftarrow R0 + R2</math></b>
x3010	1 1 1 1 0 0 0 0 0 0 1 0 0 0 0 1	<b>Print R0 (TRAP x21)</b>
x3011	1 1 1 1 0 0 0 0 0 0 1 0 0 1 0 1	<b>HALT (TRAP x25)</b>
X3012	Starting Address of File	
x3013	0 0 0 0 0 0 0 0 0 0 1 1 0 0 0 0	<b>ASCII x30 ( '0' )</b>



37

Copyright © The McGraw-Hill Companies, Inc. Permission required for reproduction or display.

## Data Path Components

- **Global bus**
  - special set of wires that carry a 16-bit signal to many components
  - inputs to the bus are “tri-state devices”, that only place a signal on the bus when they are enabled
  - only one (16-bit) signal should be enabled at any time
    - control unit decides which signal “drives” the bus
  - any number of components can read the bus
    - register only captures bus data if it is write-enabled by the control unit
- **Memory**
  - Control and data registers for memory and I/O devices
  - memory: MAR, MDR (also control signal for read/write)

CS270 - Fall 2013 - Colorado State University

38

## Data Path Components

### ● **ALU**

- Accepts inputs from register file and from sign-extended bits from IR (immediate field).
- Output goes to bus.
  - used by condition code logic, register file, memory

### ● **Register File**

- Two read addresses (SR1, SR2), one write address (DR)
- Input from bus
  - result of ALU operation or memory read
- Two 16-bit outputs
  - used by ALU, PC, memory address
  - data for store instructions passes through ALU

## Data Path Components

### ● **PC and PCMUX**

- Three inputs to PC, controlled by PCMUX
  1. PC+1 – FETCH stage
  2. Address adder – BR, JMP
  3. bus – TRAP (discussed later)

### ➤ **MAR and MARMUX**

- Two inputs to MAR, controlled by MARMUX
  1. Address adder – LD/ST, LDR/STR
  2. Zero-extended IR[7:0] -- TRAP (discussed later)

## Data Path Components

- **Condition Code Logic**
  - Looks at value on bus and generates N, Z, P signals
  - Registers set only when control unit enables them (**LD.CC**)
    - only certain instructions set the codes (**ADD, AND, NOT, LD, LDI, LDR, LEA**)
- **Control Unit – Finite State Machine**
  - On each machine cycle, changes control signals for next phase of instruction processing
    - who drives the bus? (**GatePC, GateALU, ...**)
    - which registers are write enabled? (**LD.IR, LD.REG, ...**)
    - which operation should ALU perform? (**ALUK**)
  - Logic includes decoder for opcode, etc.

## Register Transfer Notation

- **Symbolic description of what happens in a digital system**
  - higher level of abstraction than just circuit
  - Can incorporate time (sequence of events)
  - Can describe parallelism (simultaneous transfers)
- **Data path = storage elements + combinational logic**
  - **Storage elements = memories or registers**
  - **Combinational logic = operators, muxes, busses, wires**
- **Designer must specify**
  - Which transfers should take place, and in which order
  - Which control signal should be active to ensure that this transfer happens
  - Finally, use this information to design the controller FSM

## Register Transfer for Instruction Fetch

### ● Transfers

- $MAR \leftarrow PC$
- $MDR \leftarrow Mem[MAR]; PC \leftarrow PC+1$
- $IR \leftarrow MDR$

### ● Control Signals

- GatePC, LD.MAR
- MEM.EN, LD.MDR, PCMUX.Sel  $\leftarrow 10$ ; LD.PC
- Gate.MDR, LD.IR