

Chapter 5

The LC-3

Based on slides © McGraw-Hill
Additional material © 2004/2005 Lewis/Martin
Modified by Diana Palsetta

CIT 593

5-2

Instruction

Fundamental unit of work

Constituents

- **Opcode:** operation to be performed (e.g. add, and)
- **Operands:** data/locations to be used for operation
 - **Source:** location that contains the data/instruction
 - **Destination:** location that will store the result of computation
 - **Immediate:** data values not contained at a particular location

CIT 593

5-3

Instruction Set Architecture

ISA = **Programmer-visible** components & operations

- **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

All information needed to write/generate **machine language** program

CIT 593

5-2

LC-3 Overview: Memory and Registers

Memory

- Address space: **2¹⁶** locations (16-bit addresses)
- Addressability: **16 bits**

Registers

- Temporary storage (Memory access generally takes longer)
- Eight general-purpose registers: **R0 - R7** (each **16 bits wide**)
- Other registers
 - **Not directly addressable**, but used by (and affected by) instructions
 - **PC** (program counter), **condition codes**, **MAR**, **MDR**, etc.

Word Size

- **16 bits**

CIT 593

5-4

LC-3 ISA: Overview

Opcodes

- 16 opcodes ([15:12] of instruction = $2^4 = 16$ possible values)
- Types of instructions:
 - *Operate* instructions: ADD, AND, NOT
 - *Data movement* instructions: LD, LDI, LDR, LEA, ST, STR, STI
 - *Control* instructions: BR, JUMP, TRAP JSR, JSRR, RET, RTI
- Some opcodes set/clear *condition codes*, based on result
 - N = negative (<0), Z = zero (=0), P = positive (> 0)

Addressing Modes

- How is the location of an operand (data to acted upon) specified?
- Non-memory addresses: *register*, *immediate (literal)*
- Memory addresses: *base+offset*, *PC-relative*, *indirect*

Data Types

- 16-bit 2's complement integer

CIT 593


5-5

NOT (Register format)

15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0
NOT 1 0 0 1 DR SR 1 1 1 1 1 1

IR NOT R3 R5
 1 0 0 1 0 1 1 1 0 1 1 1 1 1 1 1

Convention
 source
 destination



Note: DR and SR could be the same register
 CIT 593

5-7

Operate Instructions

Only three operations

- ADD, AND, NOT

Source and destination operands are registers

- *Do not* reference memory
- ADD and AND can use “immediate” mode, (i.e., one operand is hard-wired into instruction)

Will show abstracted datapath with each instruction

- Illustrate when and where data moves to accomplish desired op.

CIT 593

5-6


ADD (Register format)

15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0
ADD 0 0 0 1 DR SR1 0 0 0 SR2

this zero means “register mode”

IR ADD R3 R5 R0
 0 0 0 1 0 1 1 1 0 1 0 0 0 0 0 0

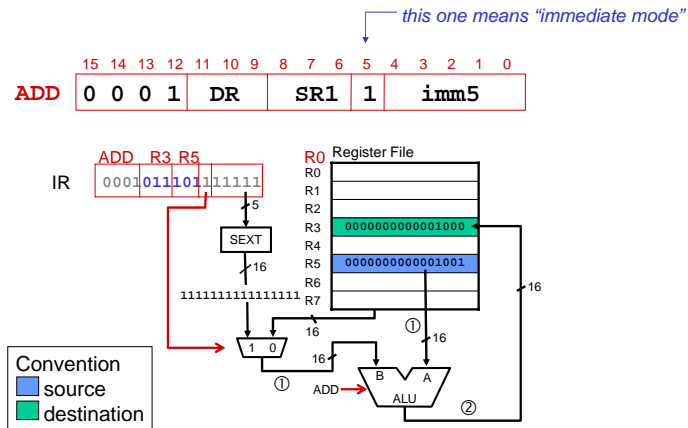
Convention
 source
 destination



CIT 593

5-8

ADD (Immediate format)



5-9

Using Operate Instructions: OR

How do we OR two numbers?

Goal

- R1 ← R2 OR R3 (no such instruction!)

Idea (Use DeMorgan's Law)

- A OR B = NOT(NOT(A) AND NOT(B))
- R4 ← NOT R2
 - R5 ← NOT R3
 - R1 ← R4 AND R5
 - R5 ← NOT R1

CIT 593

5-11

Using Operate Instructions: Subtraction

How do we subtract two numbers?

Goal

- R1 ← R2 - R3 (no such instruction!)

Idea (Use 2's complement)

- R1 ← NOT R3
- R1 ← R1 + 1
- R1 ← R2 + R1

If 2nd operand is known and small, easy

- R1 ← R2 + -3

CIT 593

5-10

Using Operate Instructions: Copying

How do we copy a number from register to register?

Goal

- R1 ← R2 (no such instruction!)

Idea (Use immediate)

- R1 ← R2 + 0

Could we use AND?

CIT 593

5-12

Using Operate Instructions: Clearing

How do we set a register to 0?

Goal

- $R1 \leftarrow 0$ (no such instruction!)

Idea

- $R1 \leftarrow R1 \text{ AND } 0$

CIT 593

5-13

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

- Take 9 bits [8:0] in instruction
- Sign extended to 16 bits
- Add it to the PC (of *next instruction*) to form address

CIT 593

5-15

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

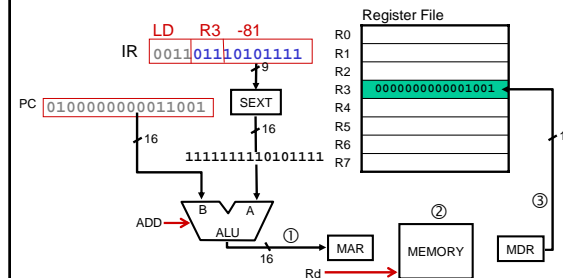
- Compute address, save in register, do not access memory
- **LEA**: immediate mode

CIT 593

5-14

LD (PC-Relative)

LD 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0
 0 0 1 0 DR PCOffset9



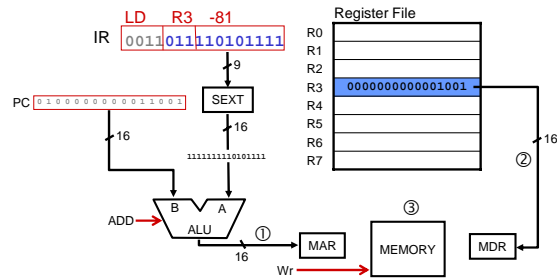
Example: LD: $R1 \leftarrow M[PC + \text{SEXT}(IR[8:0])]$

CIT 593

5-16

ST (PC-Relative)

ST 0 0 1 1 SR PCOffset9

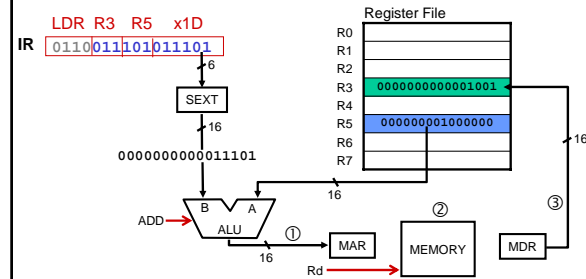


CIT 593

5-17

LDR (Base+Offset)

LDR 0 1 1 0 DR BaseR offset6



CIT 593

5-19

Base + Offset Addressing Mode

Problem

- With PC-relative mode, can only address words “near” the instruction
- What about the rest of memory?

Solution

- Use a register to generate a full 16-bit address

Idea

- 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
- *i.e.*, Instead of adding offset to PC, add it to base register

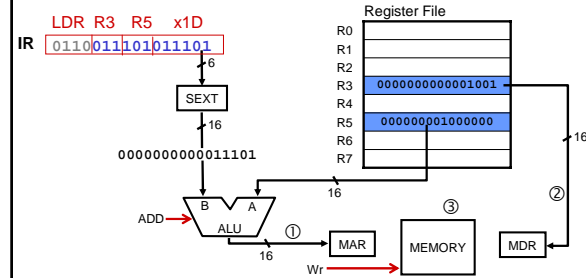
Example: LDR: R1 <- M[R2+SXT(IR[5:0])]

CIT 593

5-18

STR (Base+Offset)

STR 0 1 1 1 SR BaseR offset6



CIT 593

5-20

Indirect Addressing Mode

Another way to produce full 16-bit address

- Read address from memory location, then load/store to that address

Steps

- Address is generated from PC and PCOffset (just like PC-relative addressing)
- Then content of that address is used as address for load/store

Example: LDI: $R1 \leftarrow M[M[PC + \text{SEXT}(IR[8:0])]]$

Advantage

- Doesn't consume a register for base address
- Addresses are often stored in memory (i.e., useful)

Disadvantage

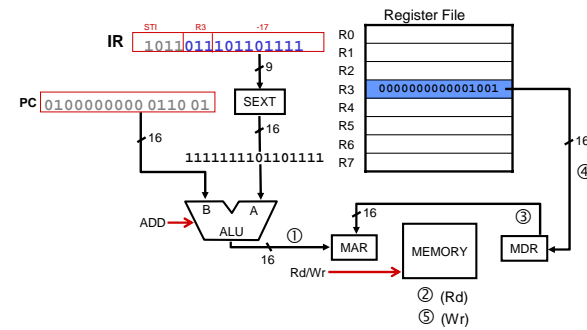
- Extra memory operation (and no offset)

CIT 593

5-21

STI (Indirect)

STI 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0
 1 0 1 1 SR PCOffset9

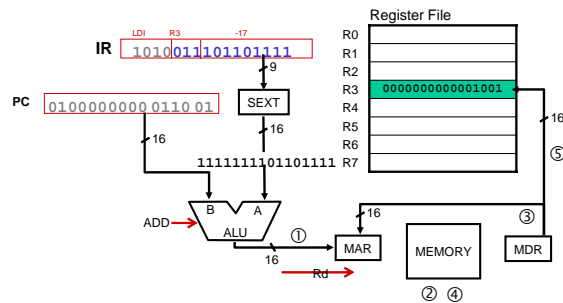


CIT 593

5-23

LDI (Indirect)

LDI 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0
 1 0 1 0 DR PCOffset9



CIT 593

5-22

Load Effective Address

Problem

- How can we compute address without also LD/ST-ing to it?

Solution

- Load Effective Address (LEA) instruction
- Used to generate address
 - Initializes a register with an address very close to the initializing instruction

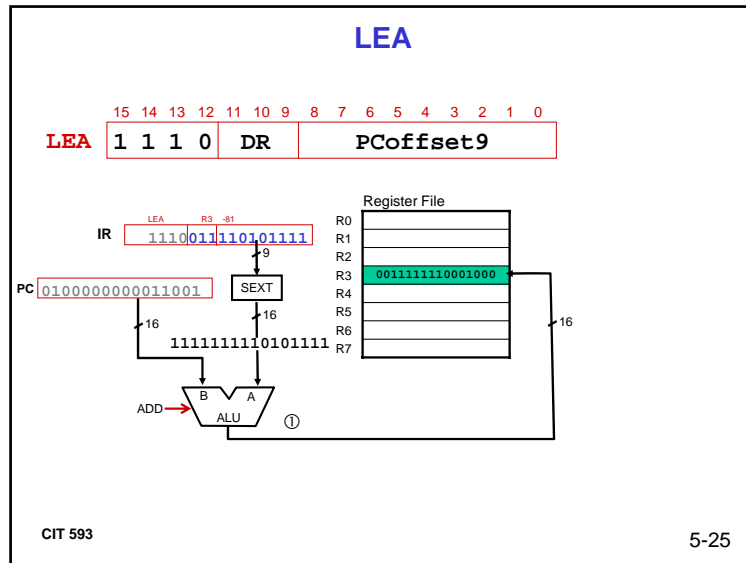
Idea

- LEA computes address just like PC-relative LD/ST
- Store address in destination register (not data at that address)
- Does not access memory

- Example: LEA: $R1 \leftarrow PC + \text{SEXT}(IR[8:0])$

CIT 593

5-24



Control Instructions

Alter the sequence of instructions

- Changing the Program Counter (PC)

Conditional Branch

- Branch *taken* if a specified condition is true
 - New PC computed relative to current PC
- Otherwise, branch *not taken*
 - PC is unchanged (i.e., points to next sequential instruction)

Unconditional Branch (or Jump)

- Always changes the PC
- Target address computed PC-relative or Base+Offset

TRAP

- Changes PC to start of OS "service routine"
- When routine is done, execution resumes after TRAP

CIT 593 5-27

Machine Language

Example

Address	opcode	Instruction	Comments
x30F6	1 1 1 0	0 0 1 1 1 1 1 1 1 0 1	$R1 \leftarrow PC-3$ (x30F4)
x30F7	0 1 0 1	0 1 0 0 1 0 1 0 0 0 0 0	$R2 \leftarrow 0$
x30F8	0 0 0 1	0 1 0 0 1 0 1 0 0 1 0 1	$R2 \leftarrow R2 + 5 = 5$
x30F9	0 1 1 1	0 1 0 0 0 1 0 0 1 1 1 0	$M[R1+14] \leftarrow R2$ $(M[x3102] \leftarrow 5)$

Example: stores the value 5 at particular location

Example: 5.10 (pg 129)
Shortened version

CIT 593 5-26

Condition Codes

LC-3 has three 1-bit **condition code** registers

N -- negative
Z -- zero
P -- positive (greater than zero)

Set/cleared by instructions that store value to register

- e.g., ADD, AND, NOT, LD, LDR, LDI, LEA, *not* ST

Exactly one will be set at all times

- Based on the last instruction that altered a register

CIT 593 5-28

Branch Instruction

Branch specifies one or more condition codes

- N, Z, or P

If the specified condition code set, the branch is taken

- PC is set to the address specified in the instruction
- Like PC-relative mode addressing, **target address** is specified as offset from current PC (PC + SEXT(IR[8:0]))
- Note: Target must be “near” branch instruction

If branch not taken i.e. the condition is not met, then next instruction (PC+1) is executed.

CIT 593

5-29

Jump Instructions

Jump is an unconditional branch (i.e., *always* taken)

So why have an special Jump instruction when we can do the same with BR?

- Because with BR we can move only in the range of the offset

Destination

- PC set to value of base register encoded in instruction
- Allows any branch target to be specified

CIT 593


5-31

BR

BR 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0
 0 0 0 0 N Z P PCoffset9

Questions

- Problems w/ this example?
- What if NZP all 0?
- What if NZP all 1?




CIT 593

5-30

JMP

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

JMP R5
 IR 1100000 101000000



CIT 593

5-32

Example: Using Branch Instructions

Goal

- Compute **sum** of 12 integers

Input


- Numbers start at x3100

Output

- Register R3 (contains **sum**)

Program

- Starts at x3000



CIT 593

5-33

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 operating system “service routine”

- Identified by 8-bit trap vector
- Execution resumes after OS code executes (more later)

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

CIT 593

5-35

Example: Summing Program

0000	BR
0001	ADD
0110	LDR
0101	AND
1110	LEA

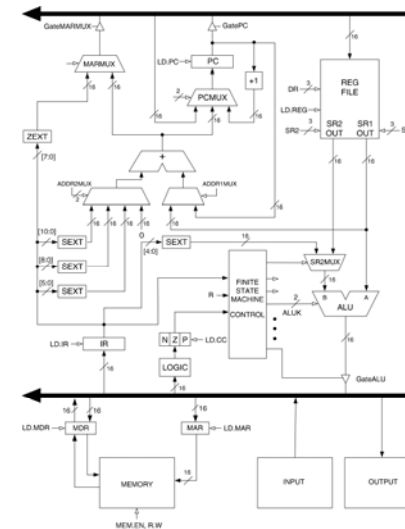
Address	Instruction	Comments
x3000	1 1 1 0 0 0 1 0 1 1 1 1 1 1 1 1	R1 ← x3001+xFF (x3100)
x3001	0 1 0 1 0 1 1 0 1 1 1 0 0 0 0 0	R3 ← 0
x3002	0 1 0 1 0 1 0 0 1 0 1 0 0 0 0 0	R2 ← 0
x3003	0 0 0 1 0 1 0 0 1 0 1 0 1 1 0 0	R2 ← 12
x3004	0 0 0 0 0 1 0 0 0 0 0 0 0 0 1 0 1	BRz x300A
x3005	0 1 1 0 1 0 0 0 0 1 0 0 0 0 0 0 0	R4 ← M[R1]
x3006	0 0 0 1 0 1 1 0 1 1 0 0 0 1 0 0	R3 ← R3+R4
x3007	0 0 0 1 0 0 1 0 0 1 1 0 0 0 0 0 1	R1 ← R1+1
x3008	0 0 0 1 0 1 0 0 1 0 1 1 1 1 1 1 1	R2 ← R2-1
x3009	0 0 0 0 1 1 1 1 1 1 1 1 1 1 0 1 0	BRnzp x3004

CIT 593

5-34

LC-3 Data Path Revisited

Filled arrow
= info to be processed.
Unfilled arrow
= control signal.



CIT 593

5-36

Data Path Components

Global bus

- Set of wires that carry 16-bit signals to many components
- Inputs to bus are controlled by triangle structure called tri-state devices (just remember the triangle for now)
 - Place signal on bus when enabled
 - Only one (16-bit) signal should be enabled at a time
 - Control unit decides which signal “drives” the bus
- Any number of components can read bus
 - Register only captures bus data if write-enabled by the control unit

Memory and I/O

- Control and data registers for memory and I/O devices
- Memory: MAR, MDR (also control signal for read/write)
- Input (keyboard): KBSR, KBDR
- Output (text display): DSR, DDR

CIT 593

5-37

Data Path Components (cont.)

PC and PCMUX

- Three inputs to PC, controlled by PCMUX
 1. Current PC plus 1 (normal operation)
 2. Adder output (BR, JMP, ...)
 3. Bus (TRAP)

MAR and MARMUX

- Some inputs to MAR, controlled by MARMUX
 1. Zero-extended IR[7:0] (used for TRAP; more later)
 2. Adder output (LD, ST, ...)

CIT 593

5-39

Data Path Components (cont.)

ALU

- Input: register file or sign-extended bits from IR (immediate field)
- Output: bus; used by...
 - Condition code registers
 - Register file
 - Memory and I/O registers

Register File

- Two read addresses, one write address (3 bits each)
- Input: 16 bits from bus
 - Result of ALU operation or memory (or I/O) read
- Outputs: two 16-bit
 - Used by ALU, PC, memory address
 - Data for store instructions passes through ALU

CIT 593

5-38

Data Path Components (cont.)

Condition Code Logic

- Looks at value on bus and generates N, Z, P signals
- Registers set only when control unit enables them
 - Only certain instructions set the codes (anything that places a value into a register: ADD, AND, NOT, LD, LDI, LDR, LEA, not ST)

Control Unit

- Decodes instruction (in IR)
- On each machine cycle, changes control signals for next phase of instruction processing
 - Who drives the bus?
 - Which registers are write enabled?
 - Which operation should ALU perform?
 - ...

CIT 593

5-40

Summary

Many instructions

- ISA: Programming-visible components and operations
- Behavior determined by opcodes and operands
 - Operate, Data, Control
- Control unit “tells” rest of system what to do (based on opcode)
- Some operations must be synthesized from given operations (e.g., subtraction, logical or, etc.)

Concepts

- Addressing modes
- Condition codes and branching/jumping

Bit-level programming bites!

CIT 593

5-41

LC-3 Instruction Summary

(inside back cover)

	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
ADD*	0001				DR			SR1		0		00				SR2
ADD*	0001				DR			SR1		1		imm5				
AND*	0101				DR			SR1		0		00				SR2
AND*	0101				DR			SR1		1		imm5				
BR	0000	n	z	p												PCoffset9
JMP	1100		000		BaseR											000000
JSR	0100		1													PCoffset11
JSRR	0100		0	00	BaseR											000000
LD*	0010				DR											PCoffset9
LDR*	1010				DR											PCoffset9
LDR*	0110				DR		BaseR									offset6
LEA*	1110				DR											PCoffset9
NOT*	1001				DR			SR								111111
RET	1100		000					111								000000
RTI	1000															0000000000
ST	0011				SR											PCoffset9
STI	1011				SR											PCoffset9
STR	0111				SR		BaseR									offset6
TRAP	1111						0000									trapvect8
reserved	1101															

CIT 593

5-43

Next Time

1-2 Examples from chapter 5 & 6

Chapter 7: Assembly Language Program

Reading: Chp6
Chp 7

CIT 593

5-42

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 a special ASCII value: **EOT (x04)**
- At the end, print the number of characters and halt (assume there will be fewer 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

CIT 593

5-44

