


Processors

Part 2

1




Processors

What are they? Besides awesome!

2

Computer Processors

- The *Central Processing Unit (CPU)* is the most complex part of a computer
- In fact, it is the computer!
- It works far different from a high-level language
- *Thousands* of processors have been developed

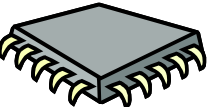


Spring 2024 Beckmann, Stein - Comp - CS5 25 3

3

Some Famous Computer Processors

- RCA 1802
- Intel 8086
- Zilog Z80
- MOS 6502
- Motorola 68000
- ARM

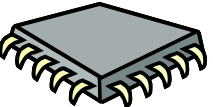


Spring 2024 Beckmann, Stein - Comp - CS5 25 4

4

Computer Processors

- Each processor functions differently
- Each is designed for a specific purpose – *form follows function*




Spring 2024 Beckmann, Stein - Comp - CS5 25 5

5

Computer Processors

- But all share some basic properties and building blocks...
- Computer hardware is divided into two "units"
 1. Control Logic Unit
 2. Execution Unit



Spring 2024 Beckmann, Stein - Comp - CS5 25 6

6

Control Logic Unit (CLU)

- *Control Logic Unit (CLU)* controls the processor
- Determines when instructions can be executed
- Controls internal operations
 - fetch & decode instructions
 - **invisible** to running programs



Spring 2024

Seacrest@Stan - CS&E - CSU 35

7

7

Execution Unit

- *Execution Unit (EU)* contains the hardware that **executes** tasks (your programs)
- Different in many processors
- Modern processors often use multiple execution units to execute instructions in parallel to improve performance

Spring 2024

Seacrest@Stan - CS&E - CSU 35

8

8

Execution Unit – The ALU

- *Arithmetic Logic Unit* is part of the Execution Unit and performs all calculations and comparisons
- Processor often contains special hardware for integer and floating point



Spring 2024

Seacrest@Stan - CS&E - CSU 35

9

9

Registers

Where the work is done

10

Registers

- In high level languages, you put active data into variables
- However, it works quite different on processors
- All computations are performed using *registers*



Spring 2024

Seacrest@Stan - CS&E - CSU 35

11

11

What – exactly – is a register?

- A *register* is a location, **on** the processor itself, that is used to store temporary data
- Think of it as a special global "variable"
- Some are accessible and usable by a programs, but **many are hidden**



Spring 2024

Seacrest@Stan - CS&E - CSU 35

12

12

What are registers used for?

- Registers are used to store anything the processor needs to keep track of
- Designed to be *fast!*
- Examples:
 - the result of calculations
 - status information
 - memory location of the running program
 - and much more...

Spring 2024

Seckman@Stan - CS68 - CS232

13

13

General Purpose Registers

- *General Purpose Registers (GPR)* don't have a specific purpose
- They are designed to be used by programs – however they are needed
- Often, you must use registers to perform calculations

Spring 2024

Seckman@Stan - CS68 - CS232

14

14

Special Registers

- There are a number of registers that are used by the Control Logic Unit and cannot be accessed by your program
- This includes registers that control how memory works, your program execution thread, and much more.

Spring 2024

Seckman@Stan - CS68 - CS232

15

15

Special Registers

- *Instruction Pointer (IP)*
 - also called *the program counter*
 - keeps track of the address of your running program
 - think it as the "line number" in your Java program – the one is being executed
 - it can be changed, but only indirectly (*using control logic – which we will cover later*)

Spring 2024

Seckman@Stan - CS68 - CS232

16

16

Special Registers

- *Status Register*
 - contains Boolean information about the processors current state
 - we will use this later, indirectly
- *Instruction Register (IR)*
 - stores the current instruction (being executed)
 - used internally and invisible to your program


Spring 2024

Seckman@Stan - CS68 - CS232

17

17

Register Files



- All the related registers are grouped into a *register file*
- Different processors access and use their register files in very different ways
- Sometimes registers are implied or hardwired

Spring 2024

Seckman@Stan - CS68 - CS232

18

18

Instructions

It's all just a bunch of bytes

19

Instructions

- You are used to writing programs in high level programming languages
- Examples:
 - C#
 - Java
 - Python
 - Visual Basic

Spring 2024 Sacramento State - CS&E - CS&E 35 20

20

High-Level Programming

- These are *third-generation languages*
- They are designed to **isolate** you from architecture of the machine
- This layer of abstraction makes programs "portable" between systems

Spring 2024 Sacramento State - CS&E - CS&E 35 21

21

Instructions

- Processors **do not** have the constructs you find in high-level languages
- Examples:
 - Blocks
 - If Statements
 - While Statements
 - ... etc

Spring 2024 Sacramento State - CS&E - CS&E 35 22

22

Instructions

- Processors can only perform a series of simple tasks
- These are called *instructions*
- *Examples:*
 - add two values together
 - copy a value
 - jump to a memory location

Spring 2024 Sacramento State - CS&E - CS&E 35 23

23

Instructions

- These instructions are used to create **all** logic needed by a program
- We will cover how to do this during the semester

Spring 2024 Sacramento State - CS&E - CS&E 35 24

24

Processor Instruction Set

- A processor's *instruction set* defines **all** the available instructions
- The instructions and their respective formats are very different for each processor

Spring 2024 Backwards Stem - CSci - CSU 25 25

25

Components of a Processor

Processor

- Execution Unit
 - Registers
 - ALU
- Control Logic Unit
 - IP
 - Status
 - IR

Spring 2024 Backwards Stem - CSci - CSU 26 26

26

The Intel x64

It was simple at first...

Spring 2024 Backwards Stem - CSci - CSU 27 27

27

The Intel x64

- The Intel x64 is the main processor used by servers, laptops, and desktops
- It has evolved continuously over a 40+ year period

Spring 2024 Backwards Stem - CSci - CSU 28 28

28

The Original x86

- First "x86" was the 8086
- Released in 1978
- Attributes:
 - 16-bit registers
 - 16 registers
 - could access of 1MB of RAM (in 64KB blocks using a special "segment" register)

Spring 2024 Backwards Stem - CSci - CSU 29 29


29

What to call the processor

- The classic term "x86" refers to the 32-bit and 16-bit processor family
- With move to 64-bit, the term "x64" is used to differentiate the newest design from the previous

Spring 2024 Backwards Stem - CSci - CSU 30 30

30




Original x86 Registers

It was simple at first...

31

Original x86 Registers

- The original x86 contained 16 registers
- 8 can be used by your programs
- The other 8 are used for memory management



Spring 2024 Sacramento State - COS - COS 35 32

32

Original x86 Registers

- The x86 processor has evolved continuously over the last 4 decades
- It first jumped to 32-bit, and then, again, to 64-bit
- The result is many of the registers have strange names

Spring 2024 Sacramento State - COS - COS 35 33

33

Original x86 Registers

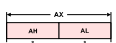
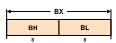
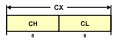

- 8 Registers can be used by your programs
 - Four General Purpose: **AX, BX, CX, DX**
 - Four pointer index: **SI, DI, BP, SP**
- The remaining 8 are restricted
 - Six segment: CS, DS, ES, FS, GS, SS
 - One instruction pointer: IP
 - One status register – used in computations

Spring 2024 Sacramento State - COS - COS 35 34

34

Original General-Purpose Registers

- However, back then (and now too) it is very useful to store 8-bit values
- So, Intel chopped 4 of the registers in half
- These registers have generic names of A, B, C, D


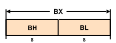
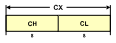






Spring 2024 Sacramento State - COS - COS 35 35

35

Original General-Purpose Registers

- The first and second byte can be used separately or used together
- Naming convention
 - high byte has the suffix "H"
 - low byte has the suffix "L"
 - for both bytes, the suffix is "X"







Spring 2024 Sacramento State - COS - COS 35 36

36

Original General-Purpose Registers

- This essentially doubled the number of registers
- So, there are:
 - four 16-bit registers or
 - eight 8-bit registers
 - ...and any combination you can think off



Spring 2024


Backwards: 8086 - 68k - CISC 35

37

37

Last the 4 Registers

- The remaining 4 registers were not cut in half
- Used for storing indexes (for arrays) and pointers
- Their purpose
 - DI – destination index
 - SI – source index
 - BP – base pointer
 - SP – stack pointer




Spring 2024

Backwards: 8086 - 68k - CISC 35

38

38

Original 16-Bit Registers



Spring 2024

Backwards: 8086 - 68k - CISC 35

39

39



Evolution to 64-Bit Registers

This is going to hurt...

40

Evolution to 32-bit

- When the x86 moved to 32-bit era, Intel expanded the registers to 32-bit
 - the 16-bit ones still exist
 - they have the prefix "e" for extended
- New instructions were added to use them




Spring 2024

Backwards: 8086 - 68k - CISC 35

41

41

Original Registers

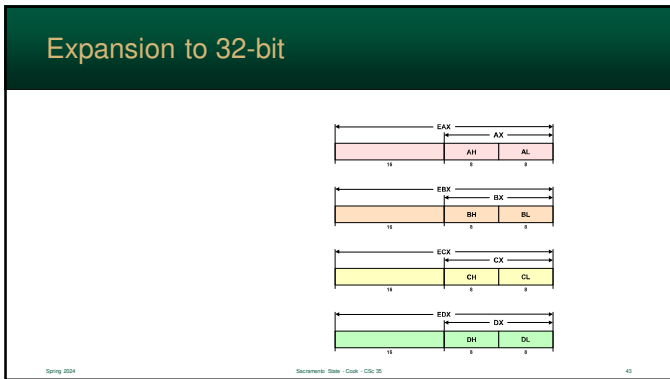


Spring 2024

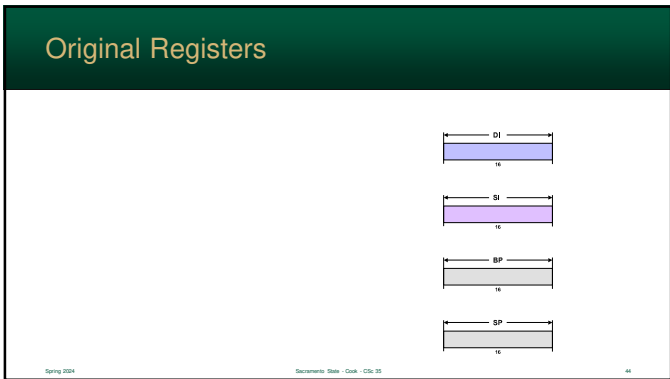
Backwards: 8086 - 68k - CISC 35

42

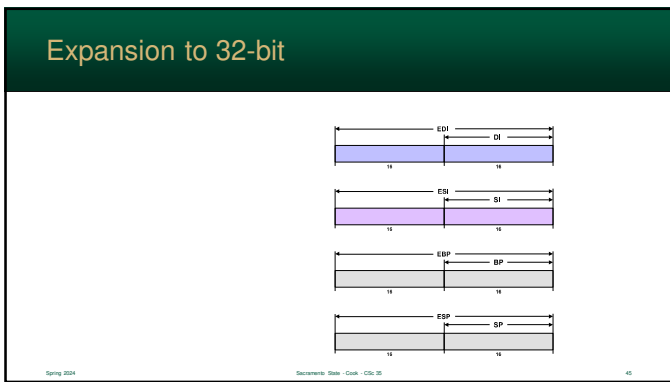
42



43



44



45

Evolution to 64-bit

- At this point, Intel had decided to **abandon** the x86 in lieu of their new Itanium Processor
- The Itanium was a radically different design and was completely incompatible
- Advanced Micro Devices (AMD), to Intel's chagrin, decided to – **once again** – extend the x86

Spring 2024 Backwards: 8086 - C++ - CSU 35 46

46

Evolution to 64-bit

- Registers were extended again
 - 64-bit registers have the prefix **"r"**
 - 8 additional registers were added
 - also, it is now possible to get 8-bit values from **all** registers (hardware is more consistent!)
- Some old, archaic, features were dropped

Spring 2024 Backwards: 8086 - C++ - CSU 35 47

47


Evolution to 64-bit

- The AMD-64 was a huge commercial **success**
- The Itanium was a commercial **failure**
- Intel, dropped the Itanium and started making 64-bit x86 using AMD's design

Spring 2024 Backwards: 8086 - C++ - CSU 35 48

48

Expansion to 64-bit




Spring 2024

Backwards Compatible - CS# - CS# 35

49

49

Expansion to 64-bit




Spring 2024

Backwards Compatible - CS# - CS# 35

50

50

Expansion to 64-bit




Spring 2024

Backwards Compatible - CS# - CS# 35

51

51

Expansion to 64-bit




Spring 2024

Backwards Compatible - CS# - CS# 35

52

52

New 64-bit Registers: R8...R15



Spring 2024

Backwards Compatible - CS# - CS# 35

53

53

64-Bit Register Table

Register	32-bit	16-bit	8-bit High	8-bit Low
rax	eax	ax	ah	al
rbx	ebx	bx	bh	bl
rcx	ecx	cx	ch	cl
rdx	edx	dx	dh	dl
rsi	esi	si		sil
rdi	edi	di		dil
rbp	ebp	bp		bpl
rsp	esp	sp		spl

Spring 2024

Backwards Compatible - CS# - CS# 35

54

54

64-Bit Register Table

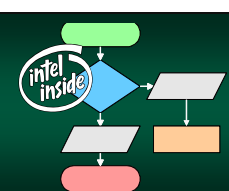
Register	32-bit	16-bit	8-bit High	8-bit Low
r8	r8d	r8w		r8b
r9	r9d	r9w		r9b
r10	r10d	r10w		r10b
r11	r11d	r11w		r11b
r12	r12d	r12w		r12b
r13	r13d	r13w		r13b
r14	r14d	r14w		r14b
r15	r15d	r15w		r15b

Spring 2024

Backwards Stem - Cosh - CS:35

55

55




Basic Intel x86 Instructions

Feel the pow-wah of the x86!

56

Basic Intel x86 Instructions

- Each x86 instruction can have up to 2 operands
- Operands in x86 instructions are very versatile
- Each operand can be either a memory address, register or an immediate value



Spring 2024

Backwards Stem - Cosh - CS:35

57

57

Types of Operands

- Registers
- Address in memory
- Register pointing to a memory address
- Immediate

Spring 2024


Backwards Stem - Cosh - CS:35

58

58

Intel x86 Instruction Limits

- There are some limitations...
- Some instructions must use an immediate
- Some instructions require a *specific* register to perform calculations



Spring 2024

Backwards Stem - Cosh - CS:35

59

59

Intel x86 Instruction Limits

- A register must always be involved
 - processors use registers for all activity
 - both operands cannot access memory at the same time
 - the processor has to have it at some point!*
- Also, obviously, the receiving field cannot be an immediate value

Spring 2024

Backwards Stem - Cosh - CS:35

60

60

Instruction: Move

- The Intel *Move Instruction* combines transfer, load and store instructions under one name
- ... well, that's something the assembler does for us – but, we'll cover that soon
- "Move" is a tad confusing – it copies data

Spring 2024

Backwards Book - Cook - CS:50

61

61

Instruction: Move

```
MOV destination, source
```



Spring 2024

Backwards Book - Cook - CS:50

62

62

Example: Move immediate

```
MOV rax, 42
```



Spring 2024

Backwards Book - Cook - CS:50

63

63

Example: Transfer

```
MOV rbx, rax
```



Spring 2024

Backwards Book - Cook - CS:50

64

64

Example: Load

```
MOV rax, total
```



Spring 2024

Backwards Book - Cook - CS:50

65

65

Example: Store

```
MOV counter, rax
```



Spring 2024

Backwards Book - Cook - CS:50

66

66

Example: "A" Register


So many options!

```
mov al, 42    #low byte
mov ah, 13    #high byte
mov ax, 1947  #both bytes
```

67

Instruction: Add & Subtract

- The Add and Subtract instructions take two operands and store the result in the first operand
- This is the same as the += and -= operators used in Visual Basic .NET, C, C++, Java, etc...



68

Instruction: Add

```
ADD target, value
```

Immediate, Register, Memory

Register, Memory

69

Example: Move register to memory

```
MOV rax, counter
ADD rax, 2
```

Move memory into rax

Same as Java
rax += 2;

70

Instruction: And & Or

- The Bitwise And & Bitwise Or instructions take two operands and stores the result in the second operand
- This is the same as the ^= and |= operators used in C, C++, Java, etc...



71

Instruction: Logical And

```
AND target, value
```

Immediate, Register, Memory

Register, Memory

72

Example: Logical Or

```
#Convert 5 to ASCII '5'  
MOV rax, 5  
OR rax, 0x30
```

0011 0000

73

Call Instruction

- The *Call Instruction* causes the processor to start running instructions at a specified memory location (a subroutine)
- Subroutines are analogous to the functions you wrote in Java
- Once it completes, execution returns from the subroutine and continues after the call

74

Call Instruction

CALL *address*

Usually a label – a constant that holds an address

75

Example: Print an integer

```
#Using the CSC35 library
```

```
MOV rax, 1846  
CALL PrintInteger
```

This name is an address

76