CMSC216: Assembly Basics and x86-64

Chris Kauffman

Last Updated: Thu Mar 13 09:29:34 AM EDT 2025

Logistics

Reading Bryant/O'Hallaron

- Skim Ch 2.7-8: Floating Point Layout
- Now Ch 3.1-7: Assembly, Arithmetic, Control
- Later Ch 3.8-11: Arrays, Structs, Floats
- Any overview guide to x86-64 assembly instructions such as Brown University's x64 Cheat Sheet

Assignments

- P2: Due Fri 14-Mar-2025
- Lab06: Assembly Coding
- HW06: Assembly Debugging

Both relevant to P3

Goals

- Floating Point Layout (20min)
- Assembly Basics
- x86-64 Overview

Announcements

Midterm Grades Posted / Grade Calculator Up

See Post 493 for info and summary statistics; overall things look good; use Grade Calculator to track your progress

Midterm Feedback Survey

Results being processed by Prof K, likely post summary stats and responses on Fri 14-Mar

Project 3 Looming

- Problem 1: Convert P2 Scale functions to assembly
- Problem 2: Puzzlebin, like Puzzlebox but no source code, use disassembled binary to trace what inputs are required to earn points

Will post before spring break BUT no expectation that students work on it over the break; \sim 14 days after break before the deadline.

Today is a Double Dot Day - Participate and earn 2 Dots

The Many Assembly Languages

- Most microprocessors are created to understand a binary machine language
- Machine Language provides means to manipulate internal memory, perform arithmetic, etc.
- The Machine Language of one processor is not understood by other processors

MOS Technology 6502

- 8-bit operations, limited addressable memory, 1 general purpose register, powered notable gaming systems in the 1980s
- Apple IIe, Atari 2600, Commodore
- Nintendo Entertainment System / Famicom

IBM Cell Microprocessor

- Developed in early 2000s, 64-bit, many cores (execution elements), many registers (32 on the PPE), large addressable space, fast multimedia performance, is a pain to program
- Playstation 3 and Blue Gene Supercomputer

Assemblers and Compilers



- Compiler: chain of tools that translate high level languages to lower ones, may perform optimizations
- Assembler: translates text description of the machine code to binary, formats for execution by processor, late compiler stage
- Consequence: The compiler can generate assembly code
- Generated assembly is a pain to read but is often quite fast
- Consequence: A compiler on an Intel chip can generate assembly code for a different processor, cross compiling

Our focus: The x86-64 Assembly Language

- ► x86-64 Targets Intel/AMD chips with 64-bit word size Reminder: 64-bit "word size" ≈ size of pointers/addresses
- Lineage of x86 family
 - 1970s: 16-bit systems like Intel 8086
 - 1990s: IA32 (Intel 32-bit systems like 80386 and 80486)
 - 2000s: x86-64 (64-bit extension by AMD)

▶ x86-64 is backwards compatibility, consequently much cruft

- Can run compiled code from the 70's / 80's on modern processors without much trouble BUT means 50-year-old instructions must be preserved
- x86-64 is not the assembly language you would design from scratch today, it's the assembly you have to code against
- RISC-V is a new assembly language that is "clean" as it has no history to support (and few CPUs run it)
- Warning: Lots of information available on the web for Intel assembly programming BUT some of it is dated, IA32 info which may not work on 64-bit systems

x86-64 Assembly Language Syntax(es)

- Different assemblers understand different syntaxes for the same assembly language
- GCC use the GNU Assembler (GAS, command 'as file.s')
- GAS and Textbook favor AT&T syntax so we will too
- NASM assembler favors Intel, may see this online

AT&T Synta	ax (Our Focus)	Intel Syntax		
multstore: pushq movq call movq popq ret	%rbx %rdx, %rbx mult2@PLT %rax, (%rbx) %rbx	multstore: push mov call mov pop ret	rbx rbx, rdx mult20PLT QWORD PTR rbx	[rbx], rax

- Use of % to indicate registers
- Use of q/1/w/b to indicate 64 / 32 / 16 / 8-bit operands

	Register	names	are	bare
--	----------	-------	-----	------

 Use of QWORD etc. to indicate operand size

Generating Assembly from C Code

- gcc -S file.c will stop compilation at assembly generation
- Leaves assembly code in file.s
 - file.s and file.S conventionally assembly code though sometimes file.asm is used
- By default, compiler generates code that is often difficult for humans to interpret, may include re-arrangements, "conservative" compatibility assembly, etc. increasing size of assembly considerably
- gcc -Og file.c: optimize for debugging, generally makes it easier to read generated assembly, aligns somewhat more closely to C code

Example of Generating Assembly from C

```
# show C file to be translated
>> cat exchange.c
// exchange.c: sample C function
// to compile to assembly
long exchange(long *xp, long y){
                                       # function to translate
 long x = *xp;
                                          # involves pointer deref
 *xp = y;
 return x;
>> gcc -Og -S exchange.c
                                          # Compile to show assembly
                                          # -Og: debugging level optimization
                                          # -S: only output assembly
>> cat exchange.s
                                          # show assembly output
        .file "exchange.c"
        text
        .globl exchange
        .type exchange, @function
                                          # beginning of exchange function
exchange:
.LFBO:
        .cfi startproc
       movq (%rdi), %rax
                                          # pointer derefs in assembly
        movq %rsi, (%rdi)
                                          # uses registers
        ret.
        .cfi_endproc
.LFEO:
        .size exchange, .-exchange
        .ident "GCC: (GNU) 11.1.0"
        .section .note.GNU-stack,"", Oprogbits
```

gcc -Og -S mstore.c

```
> cat mstore.c
                                           # show a C file
long mult2(long a, long b);
void multstore(long x, long y, long *dest){
  long t = mult2(x, y);
  *dest = t:
> gcc -Og -S mstore.c
                                           # Compile to show assembly
                                           # -Og: debugging level optimization
                                           # -S: only output assembly
> cat mstore.s
                                           # show assembly output
        .file "mstore.c"
        text
        .globl multstore
                                           # function symbol for linking
        .type
               multstore. @function
multstore:
                                           # beginning of mulstore function
.LFBO:
                                           # assembler directives
        .cfi_startproc
        pushq %rbx
                                           # assembly instruction
        .cfi def cfa offset 16
                                           # directives
        .cfi offset 3. -16
        movq %rdx, %rbx
                                           # assembly instructions
        call mult20PLT
                                           # function call
       movq %rax, (%rbx)
       popq
              %rbx
        .cfi_def_cfa_offset 8
                                           # function return
        ret
        .cfi endproc
```

Every Programming Language

Look for the following as it should almost always be there

- Comments
- □ Statements/Expressions
- Variable Types
- Assignment
- Basic Input/Output
- Function Declarations
- □ Conditionals (if-else)
- □ Iteration (loops)
- 🗆 Aggregate data (arrays, structs, objects, etc)
- Library System

Exercise: Examine col_simple_asm.s

Take a simple sample problem to demonstrate assembly:

Computes Collatz Sequence starting at n=10: if n is ODD n=n*3+1; else n=n/2. Return the number of steps to converge to 1 as the **return code** from main()

The following codes solve this problem

Code	Notes
col_simple_asm.s	Hand-coded assembly for obvious algorithm
	Straight-forward reading
col_unsigned.c	Unsigned C version
	Generated assembly is reasonably readable
col_signed.c	Signed C vesion
	Generated assembly is interesting

- Kauffman will Compile/Run code
- Students should study the code and predict what lines do
- Illustrate tricks associated with gdb and assembly

Exercise: col_simple_asm.s

```
1 ### Compute Collatz sequence starting at 10 in assembly.
 2 .section .text
 3 .globl main
 4 main:
                   $0. %r8d
                                  # int steps = 0;
 5
           movl
           movl
                   $10, %ecx
                                   # int n = 10;
 6
 7
   .LOOP:
                   $1, %ecx
                                   # while(n > 1){ // immediate must be first
 8
           cmpl
 9
           jle
                   .END
                                   #
                                       n <= 1 exit loop
10
           movl
                   $2, %esi
                                     divisor in esi
                                    #
                   %ecx,%eax
                                   # prep for division: must use edx:eax
11
           movl
12
           cqto
                                    #
                                     extend sign from eax to edx
           idivl
                   %esi
                                       divide edx:eax by esi
13
                                   #
                                   #
                                       eax has quotient. edx remainder
14
                                     if(n % 2 == 1) {
15
           cmpl
                   $1,%edx
                                   #
                   .EVEN
                                         not equal, go to even case
16
           jne
                                    #
17
18
           imull
                   $3, %ecx
                                   #
                                        n = n * 3
           incl
                   %ecx
                                   #
                                         n = n + 1 OR n++
19
                                       }
                   .UPDATE
                                    #
20
           jmp
   .EVEN:
                                   #
                                       else{
21
           sarl
                   $1.%ecx
                                   #
                                         n = n / 2: via right shift
22
                                   #
                                        3
23
   . UPDATE:
                   %r8d
24
           incl
                                    #
                                        steps++:
                                   # }
25
           jmp
                   .LOOP
26
   .END:
                   %r8d. %eax
                                   # r8d is steps, move to eax for return value
27
           movl
28
           ret
29
```

Answers: x86-64 Assembly Basics for AT&T Syntax

- Comments are one-liners starting with #
- Statements: each line does ONE thing, frequently text representation of an assembly instruction

movq %rdx, %rbx # move rdx register to rbx

Assembler directives and labels are also possible:

.global multstore	#	notify linker	of	loca	ation	mult	store
multstore:	#	label beginnin	ıg	of mu	iltsto	ore s	ection
<mark>blah</mark> blah blah	#	instructions i	n	this	this	sect	ion

- Variables: mainly registers, also memory ref'd by registers maybe some named global locations
- Assignment: instructions like movX that put bits into registers and memory
- Conditionals/Iteration: assembly instructions that jump to code locations
- Functions: code locations that are **labeled** and global
- Aggregate data: none, use the stack/multiple registers
- Library System: link to other code

So what are these Registers?

- Memory locations directly wired to the CPU
- Usually very fast to access, faster than main memory
- Most instructions involve registers, access or change reg val

Example: Adding Together Integers

- Ensure registers have desired values in them
- Issue an addX instruction involving the two registers
- Result will be stored in a register

```
addl %eax, %ebx
# add ints in eax and ebx, store result in ebx
addq %rcx, %rdx
# add longs in rcx and rdx, store result in rdx
```

Note instruction and register names indicate whether 32-bit int or 64-bit long are being added

x86-64 "General Purpose" Registers

Many "general purpose" registers have special purposes and conventions associated such as

- Return Value:
 %rax / %eax / %ax
- Function Args 1 to 6: %rdi, %rsi, %rdx, %rcx, %r8, %r9
- Stack Pointer (top of stack): %rsp
- Old Code Base Pointer: %rbp, historically start of current stack frame but is not used that way in modern codes

Note: There are also Special Registers like %rip and %eflags which we will discuss later.

64-bit	32-bit	16-bit	8-bit	Notes	
%rax	%eax	%ax	%al	Return Val	
%rbx	%ebx	%bx	%bl		
%rcx	%ecx	%CX	%cl	Arg 4	
%rdx	%edx	%dx	%dl	Arg 3	
%rsi	%esi	%si	%sil	Arg 2	
%rdi	%edi	%di	%dil	Arg 1	
%rsp	%esp	%sp	%spl	Stack Ptr	
%rbp	%ebp	%bp	%bpl	Base Ptr?	
%r8	%r8d	%r8w	%r8b	Arg 5	
%r9	%r9d	%r9w	%r9b	Arg 6	
%r10	%r10d	%r10w	%r10b		
%r11	%r11d	%r11w	%r11b		
%r12	%r12d	%r12w	%r12b		
%r13	%r13d	%r13w	%r13b		
%r14	%r14d	%r14w	%r14b		
%r15	%r15d	%r15w	%r15b		
Caller	Save:	Restore	after cal	ling func	
Callee	Save:	Restore before returning			

Register Naming Conventions

- AT&T syntax identifies registers with prefix %
- Naming convention is a historical artifact
- Originally 16-bit architectures in x86 had
 - General registers ax, bx, cx, dx,
 - Special Registers si,di,sp,bp
- Extended to 32-bit: eax,ebx,...,esi,edi,...
- Grew again to 64-bit: rax,rbx,...,rsi,rdi,...
- Added Eight 64-bit regs r8,r9,...,r14,r15 with 32-bit portion r8d,r9d,..., 16-bit r8w,r9w..., etc.
- Instructions must match registers sizes:

addw %ax, %bx # word (16-bit) addl %eax, %ebx # long word (32-bit) addq %rax, %rbx # quad-word (64-bit)

When hand-coding assembly, easy to mess this up, assembler will error out

Hello World in x86-64 Assembly : Not that Easy

- Non-trivial in assembly because output is involved
 - Try writing helloworld.c without printf()
- Output is the business of the operating system, always a request to the almighty OS to put something somewhere
 - Library call: printf("hello"); mangles some bits but eventually results with a ...
 - System call: Unix system call directly implemented in the OS kernel, puts bytes into files / onto screen as in write(1, buf, 5); // file 1 is screen output

This gives us several options for hello world in assembly:

- hello_printf64.s: via calling printf() which means the C standard library must be (painfully) linked
- hello64.s via direct system write() call which means no external libraries are needed: OS knows how to write to files/screen. Use the 64-bit Linux calling convention.
- 3. hello32.s via direct system call using the older 32 bit Linux calling convention which "traps" to the operating system.

(Optional): The OS Privilege: System Calls

- Most interactions with the outside world happen via Operating System Calls (or just "system calls")
- User programs indicate what service they want performed by the OS via making system calls
- System Calls differ for each language/OS combination
 - x86-64 Linux: set %rax to system call number, set other args in registers, issue syscall
 - IA32 Linux: set %eax to system call number, set other args in registers, issue an interrupt
 - C Code on Unix: make system calls via write(), read() and others (studied in CSCI 4061)
 - Tables of Linux System Call Numbers
 - 64-bit (335 calls)
 - ▶ 32-bit (190 calls)
 - Mac OS X: very similar to the above (it's a Unix)
 - Windows: use OS wrapper functions
- OS executes priveleged code that can manipulate any part of memory, touch internal data structures corresponding to files, do other fun stuff discussed in CSCI 4061 / 5103

Basic Instruction Classes

Remember: Goal is to understand assembly as a *target* for higher languages, not become expert "assemblists"

- Means we won't hit all 4,834 pages of the Intel x86-64 Manual
- Brown University's x64 Cheat Sheet has a good overview
- x86 Assembly Guide from Yale is also good but is limited to 32-bit coverage

Kind	Assembly Instructions
Fundamentals	
- Memory Movement	mov
- Stack manipulation	push,pop
- Addressing modes	(%eax),12(%eax,%ebx)
Arithmetic/Logic	
- Arithmetic	add,sub,mul,div,lea
- Bitwise Logical	and,or,xor,not
- Bitwise Shifts	sal,sar,shr
Control Flow	
- Compare / Test	cmp,test
- Set on result	set
- Jumps (Un)Conditional	jmp,je,jne,jl,jg,
- Conditional Movement	cmove, cmovg,
Procedure Calls	
- Stack manipulation	push,pop
- Call/Return	call,ret
- System Calls	syscall
Floating Point Ops	
- FP Reg Movement	vmov
- Conversions	vcvts
- Arithmetic	vadd,vsub,vmul,vdiv
- Extras	vmins,vmaxs,sqrts

Data Movement: movX instruction

movX SOURCE, DEST

move/copy source value to dest

Overview

- Moves data...
 - Reg to Reg
 - Mem to Reg
 - Reg to Mem
 - Imm to . . .
- Reg: register
- Mem: main memory
- Imm: "immediate" value (constant) specified like
 - \$21 : decimal
 - \$0x2f9a : hexadecimal
 - NOT 1234 (mem adder)
- More info on operands next

Examples

64-bit quadword moves movq \$4, %rbx # rbx = 4; movq %rbx,%rax # rax = rbx; movq \$10, (%rcx) # *rcx = 10;

32-bit longword moves
movl \$4, %ebx # ebx = 4;
movl %ebx,%eax # eax = ebx;
movl \$10, (%rcx) # *rcx = 10;
Note variations

- movq for 64-bit (8-byte)
- movl for 32-bit (4-byte)
- movw for 16-bit (2-byte)
- movb for 8-bit (1-byte)

Operands and Addressing Modes

In many instructions like movX, operands can have a variety of forms called **addressing modes**, may include constants and memory addresses

Style	Address Mode	C-like	Notes
\$21	immediate	21	value of constant like 21
\$0xD2			or $0xD2 = 210$
%rax (%rax) 8(%rax) 4(%rdx)	register indirect displaced	rax *rax *(rax+2) rdx->field	to/from register contents reg holds memory address, deref base plus constant offset, often used for strcut field derefs
(%rax,%rbx)	indexed	*(rax+rbx) char_arr[rbx]	base plus offset in given reg actual value of rbx is used, NOT multiplied by sizeof()
(%rax,%rbx,4) (%rax,%rbx,8)	scaled index	rax[rbx] rax[rbx]	<pre>like array access with sizeof()=4 "" with sizeof()=8</pre>
1024	absolute		Absolute address #1024 Rarely used

Exercise: Show movX Instruction Execution

```
Code movX exercise.s
                              Registers/Memory
movl $16, %eax
                              TNTTTAL.
movl $20, %ebx
                               movq $24, %rbx
                               REG | %rax |
## POS A
                                    | %rbx |
                                    | %rcx | #1024
movl %eax,%ebx
                                    | %rdx | #1032
movq %rcx,%rax
                                ----+-----+-----
## POS B
                               MEM | #1024 |
                                    | #1032 | 25 |
movq $45,(%rdx)
                                            15 |
                                     #1040 |
movl $55,16(%rdx)
                                             5
                                    #1048 |
## POS C
                               ----+----+-------
                              Lookup...
movq $65, (%rcx, %rbx)
                              May need to look up addressing
movq $3,%rbx
```

POS D

movq \$75, (%rcx,%rbx,8) conventions for things like...

> movX %y,%x # reg y to reg x movX \$5,(%x) # 5 to address in %x

0

0

35 I

Answers Part 1/2: movX Instruction Execution



#WARNING!: On 64-bit systems, ALWAYS use a 64-bit reg name like %rdx and movq to copy memory addresses; using smaller name like %edx will miss half the memory addressing leading to major memory problems

Answers Part 2/2: movX Instruction Execution

		movq \$65,(%rcx,%rbx)			
	movq \$45,(%rdx)	#1024+16 = #			
movl %eax,%ebx	#1032	movq \$3,%rbx			
movq %rcx,%rax #!	movq \$55,16(%rdx)	movq \$75,(%rcx,%rbx,8)			
-	16+#1032=#1048	#1024 + 3*8			
## POS B	## POS C	## POS D			
REG VALUE	REG VALUE	REG VALUE			
%rax #1024	%rax #1024	%rax #1024			
%rbx 16	%rbx 16	%rbx 3			
%rcx #1024	%rcx #1024	%rcx #1024			
%rdx #1032	%rdx #1032	%rdx #1032			
MEM VALUE	MEM VALUE	MEM VALUE			
#1024 35	#1024 35	#1024 35			
#1032 25	#1032 45	#1032 45			
#1040 15	#1040 15	#1040 65			
#1048 5	#1048 55	#1048 75			

#1024+16 = #1040

 $#1024 + 3 \times 8 = #1048$

gdb Assembly: Examining Memory

gdb commands print and x allow one to print/examine memory memory of interest. Try on movX_exercises.s

```
# TUI mode
(gdb) tui enable
(gdb) layout asm
                         # assembly mode
(gdb) layout reg
                       # show registers
(gdb) stepi
                           # step forward by single Instruction
(gdb) print $rax
                    # print register rax
(gdb) print *($rdx)
                          # print memory pointed to by rdx
(gdb) print (char *) $rdx # print as a string (null terminated)
(gdb) x $r8
                           # examine memory at address in r8
(gdb) x/3d $r8
                           # same but print as 3 4-byte decimals
(gdb) x/6g $r8
                           # same but print as 6 8-byte decimals
(gdb) x/s $r8
                           # print as a string (null terminated)
(gdb) print *((int*) $rsp) # print top int on stack (4 bytes)
(gdb) x/4d $rsp
                           # print top 4 stack vars as ints
(gdb) x/4x $rsp
                           # print top 4 stack vars as ints in hex
```

Many of these tricks are needed to debug assembly.

Register Size and Movement

- Recall %rax is 64-bit register, %eax is lower 32 bits of it
- Data movement involving small registers may NOT overwrite higher bits in extended register
- Moving data to low 32-bit regs automatically zeros high 32-bits movabsq \$0x1122334455667788, %rax # 8 bytes to %rax movl \$0xAABBCCDD, %eax # 4 bytes to %eax ## %rax is now 0x0000000AABBCCDD
- Moving data to other small regs DOES NOT ALTER high bits movabsq \$0x1122334455667788, %rax # 8 bytes to %rax movw \$0xAABB, %ax # 2 bytes to %ax ## %rax is now 0x112233445566AABB

Exercise: movX differences in Main Memory

Instr	# bytes	
movb	1 byte	
movw	2 bytes	
movl	4 bytes	
movq	8 bytes	

Show the result of each of the following copies to main memory in sequence.

movl	%eax,	(%rsi)	#1
movq	%rax,	(%rsi)	#2
movb	%cl,	(%rsi)	#3
movw	%cx,	2 (%rs i)	#4
movl	%ecx,	4(%rsi)	#5
movw	4 (%rs :	i), %ax	#6

INITIAL	
INITIAL REG rax rcx rsi MEM #1024 #1025	0x000000000000000000000000000000000000
#1026 #1027	0x22 0x33
#1028	0x44
#1029	0x55
#1030	0x66
#1031	0x77
#1032	0x88
#1033 	0x99

Answers: movX to Main Memory 1/2

+			movl	%eax,	(%rsi)	#1 4	bytes	rax ->	#1024	
REG		- I	movq	%rax,	(%rsi)	#2 8	bytes	rax ->	#1024	
rax 0x	0000000DDC	CBBAA	movb	%cl,	(%rsi)	#3 1	byte	rcx ->	#1024	
rcx 0x	000000000000000	OFFEE	movw	%cx,	2(%rsi)	#4 2	bytes	rcx ->	#1026	
rsi		#1024	movl	%ecx,	4(%rsi)	#5 4	bytes	rcx ->	#1028	
+			movw	4(%rs	i), %ax	#6 2	bytes	#1024 ·	-> rax	
	#	1		#2			#3	3		
INITIAL	m	ovl %eax	,(%rsi)	mo	vq %rax,	(%rsi)) ma	ovb %cl	,(%rsi)	
+-		+		-	+		-		+	L
MEM		MEM	- I	- I	MEM		I I	MEM	I	I
#1024	0x00	#1024	OxAA	1	#1024 (OxAA		#1024	0xEE	L
#1025	0x11	#1025	OxBB	1	#1025 (OxBB		#1025	OxBB	L
#1026	0x22	#1026	0xCC	1	#1026 (0xCC		#1026	0xCC	L
#1027	0x33	#1027	OxDD	1	#1027 (OxDD		#1027	OxDD	L
#1028	0x44	#1028	0x44	1	#1028 (0x00		#1028	0x00	L
#1029	0x55	#1029	0x55	1	#1029 (0x00		#1029	0x00	L
#1030	0x66	#1030	0x66	1	#1030 (0x00		#1030	0x00	L
#1031	0x77	#1031	0x77	- I -	#1031 (0x00		#1031	0x00	I
#1032	0x88	#1032	0x88	- I -	#1032 (0x88		#1032	0x88	I
#1033	0x99	#1033	0x99	1	#1033 (0x99		#1033	0x99	L
+-		+-		-	+		-		+	L

Answers: movX to Main Memory 2/2

	+			movl	%eax,	(%rsi)	#1 4	bytes	rax ->	#10)24	
	REG		1	movq	%rax,	(%rsi)	#2 8	bytes	rax ->	#10)24	
	rax	0x00000	000DDCCBBAA	movb	%cl,	(%rsi)	#3 1	byte	rcx ->	#10)24	
	rcx	0x00000	0000000FFEE	movw	%cx,	2(%rsi)	#4 2	bytes	rcx ->	#10	26	
	rsi		#1024	movl	%ecx,	4(%rsi)	#5 4	bytes	rcx ->	#10)28	
	+			movw	4(%rs:	i), %ax	#6 2	bytes	#1024 ·	-> 1	ax	
4	#3		#4		#!	5		1	#6			
I	novb %c	l,(%rsi)	movw %cx	,2(%rsi)	m	ovl %ecx	,4(%rs	si) r	novw 4(?	(rsz	r),%a	x
		-+		+	ŀ	+		-		-+		L
	MEM	1	MEM	I I	- I	MEM		1	MEM	I.		L
	#1024	OxEE	#1024	0xEE	- I	#1024	0xEE	1	#1024	10)xEE	L
	#1025	OxBB	#1025	OxBB	- I	#1025	OxBB	1	#1025	10)xBB	L
	#1026	0xCC	#1026	0xEE	- I	#1026	OxEE	1	#1026	10)xEE	L
	#1027	OxDD	#1027	0xFF	- I	#1027	0xFF	1	#1027	10)xFF	L
	#1028	0x00	#1028	0x00	- I	#1028	0xEE	1	#1028	10)xEE	<
	#1029	0x00	#1029	0x00	- I	#1029	0xFF	1	#1029	10)xFF	<
	#1030	0x00	#1030	0x00	- I	#1030	0x00	1	#1030	10	00x0	L
	#1031	0x00	#1031	0x00	- I	#1031	0x00	1	#1031	10	00x0	L
	#1032	0x88	#1032	0x88	- I	#1032	0x88	1	#1032	10)x88	L
	#1033	0x99	#1033	0x99	- I	#1033	0x99	I	#1033	10)x99	L
		-+		+	1.	+		-		-+		1

| rax | 0x0000000DDCCFFEE |

addX : A Quintessential ALU Instruction

addX B, A # A = A+B

OPERANDS:

addX %reg, %reg addX (%mem),%reg addX %reg, (%mem) addX \$con, %reg addX \$con, (%mem)

No mem+mem or con+con

EXAMPLES:

addq	%rdx,	%rcx	#	rcx	=	rcx	+	rċ	lx
addl	%eax,	%ebx	#	ebx	=	ebx	+	ea	аx
addq	\$42,	%rdx	#	rdx	=	rdx	+	42	2
addl	(%rsi)	,%edi	#	edi	=	edi	+	*1	rsi
addw	%ax,	(%rbx)	#	*rbx	=	*rb	хc	+	ax
addq	\$55,	(%rbx)	#	*rbx	=	*rb	хc	+	55

addl (%rsi,%rax,4),%edi # edi = edi+rsi[rax] (int)

- Addition represents most 2-operand ALU instructions well
- Second operand A is modified by first operand B, No change to B
- Variety of register, memory, constant combinations honored
- addX has variants for each register size: addq, addl, addw, addb

Optional Exercise: Addition

Show the results of the following addX/movX ops at each of the specified positions

```
addq $1,%rcx # con + reg
addg %rbx,%rax # reg + reg
## POS A
addq (%rdx),%rcx # mem + reg
addq %rbx,(%rdx) # reg + mem
addg 3, (%rdx) # con + mem
## POS B
addl $1,(%r8,%r9,4) # con + mem
addl $1,%r9d
                     # con + reg
addl %eax,(%r8,%r9,4) # reg + mem
addl $1,%r9d
                      # con + reg
addl (%r8,%r9,4),%eax
                      # mem + reg
## POS C
```

INITIAL							
	+						
REGS							
%rax	15						
%rbx	20						
%rcx	25						
%rdx	#1024						
%r8	#2048						
1 % ~ 9							
1 /01 3	1 01						
	+						
 MEM	+ 						
 MEM #1024	+ 100						
 MEM #1024 	0 + 100 						
 MEM #1024 #2048	0 100 200						
 MEM #1024 #2048 #2052	0 100 200 300						
 MEM #1024 #2048 #2052 #2056	+ 100 200 300 400						

Answers: Addition

I	NITIAL				P	JS A				Ρ	OS B				P	DS C			
Ŀ		+-		L	1-		+		-1	I		+		٠I	ŀ		+-		-1
L	REG	L		L	L	REG	I		L	I	REG	I		L	Т	REG	T		1
L	%rax	L	15	L	L	%rax	I	35	T	I	%rax	I	35	L	Т	%rax	T	435	I
L	%rbx	L	20	L	L	%rbx	I	20	T	I	%rbx	I	20	L	Т	%rbx	T	20	I
L	%rcx	L	25	L	L	%rcx	I	26	T	I	%rcx	I	126	L	Т	%rcx	T	126	I
L	%rdx	L	#1024	L	L	%rdx	I	#1024	T	I	%rdx	I	#1024	L	Т	%rdx	T	#1024	I
L	%r8	L	#2048	L	L	%r8	L	#2048	T	I	%r8	I	#2048	L	Т	%r8	T	#2048	I
L	%r9	L	0	L	L	%r9	L	0	L	I	%r9	I	0	L	T	%r9	T	2	I
Ŀ		+-		1	1-		+		•			+		٠I	ŀ		+•		-1
L	MEM	L		L	L	MEM	L		L		MEM	I		L		MEM	Ι		
L	#1024	L	100	L	L	#1024	L	100	L	I	#1024	I	123	L	T	#1024	T	123	I
L		L		L	L		L		L	I		I		L	T		T		I
L	#2048	L	200	L	L	#2048	L	200	L	I	#2048	I	200	L	T	#2048	T	201	I
L	#2052	L	300	L	L	#2052	L	300	T	I	#2052	I	300	L	Т	#2052	T	335	I
l	#2056	L	400	L	L	#2056	I	400	L	I	#2056	I	400	L	Т	#2056	T	400	I
Ŀ		+-		L	1-		+		•			+		·١	ŀ		++		-1

addq	\$1,%rcx	addq	(%rdx),%rcx	addl	\$1,(%r8,%r9,4)
addq	%rbx,%rax	addq	%rbx,(%rdx)	addl	\$1,%r9d
		addq	\$3,(%rdx)	addl	%eax,(%r8,%r9,4)
				addl	\$1,%r9d
				addl	(%r8,%r9,4),%eax

33

The Other ALU Instructions

- Most ALU instructions follow the same patter as addX: two operands, second gets changed.
- Some one operand instructions as well.

Instruction	Name	Effect	Notes
addX B, A	Add	A = A + B	Two Operand Instructions
subX B, A	Subtract	A = A - B	
imulX B, A	Multiply	A = A * B	Has a limited 3-arg variant
andX B, A	And	A = A & B	
orX B, A	Or	A = A B	
xorX B, A	Xor	$A = A \cap B$	
salX B, A	Shift Left	$A = A \ll B$	B is constant or %cl reg
shlX B, A		A = A << B	
sarX B, A	Shift Right	A = A >> B	Arithmetic: Sign carry
shrX B, A		A = A >> B	Logical: Zero carry
incX A	Increment	A = A + 1	One Operand Instructions
decX A	Decrement	A = A - 1	
negX A	Negate	A = -A	
notX A	Complement	A = ~A	

leaX: Load Effective Address

- Memory addresses must often be loaded into registers
- Often done with a leaX, usually leaq in 64-bit platforms
- Sort of like "address-of" op & in C but a bit more general

INITIAL		## leaX_example
	+	movq 8(%rdx),%
REG	VAL	<pre>leaq 8(%rdx),%</pre>
rax	0	movl (%rsi,%rc
rcx	2	leaq (%rsi,%rc
rdx rsi 	#1024 #2048 +	Compiler somet as it is usually f
MEM #1024 #1032 #2048 #2052 #2056 	15 25 200 300 400	<pre># Odd Collatz u #READABLE with imul \$3,%eax addl \$1,%eax # eax = eax*3 + # 3-4 cycles</pre>
1	•	

```
# leaX_examples.s:
ovq 8(%rdx),%rax # rax = *(rdx+1) = 25
eaq 8(%rdx),%rax # rax = rdx+1 = #1032
ovl (%rsi,%rcx,4),%eax # rax = rsi[rcx] = 400
eaq (%rsi,%rcx,4),%rax # rax = &(rsi[rcx]) = #2056
```

Compiler sometimes uses leaX for multiplication as it is usually faster than imulX but less readable.

```
# Odd Collatz update n = 3*n+1
#READABLE with imulX  #OPTIMIZED with leaX:
imul $3,%eax
addl $1,%eax
# eax = eax*3 + 1  # eax = eax + 2*eax + 1,
# 3-4 cycles  # 1 cycle
```



Division: It's a Pain (1/2)

- idivX operation has some special rules
- Dividend must be in the rax / eax / ax register
- Sign extend to rdx / edx / dx register with cqto
- idivX takes one register argument which is the divisor
- At completion

answer in eax, return ret

Compiler avoids division whenever possible: compile col_unsigned.c and col_signed.c to see some tricks.

Division: It's a Pain (2/2)

When performing division on 8-bit or 16-bit quantities, use instructions to sign extend small reg to all rax register

```
### division with 16-bit shorts from division.s
movg $0,%rax
              # set rax to all 0's
movq $0,%rdx
                 # set rdx to all 0's
                  \# rax = 0x0000000 00000000
                  \# rdx = 0x0000000 0000000
movw $-17, %ax
                 # set ax to short -17
                  \# rax = 0x0000000 0000FFEF
                  \# rdx = 0x0000000 00000000
cwtl
                  # "convert word to long" sign extend ax to eax
                  \# rax = 0x0000000 FFFFFFFF
                  \# rdx = 0x0000000 00000000
                  # "convert long to quad" sign extend eax to rax
cltq
                  \# rdx = 0x0000000 00000000
                  # sign extend rax to rdx
cqto
                  movq $3, %rcx
                  # set rcx to long 3
idivq %rcx
                  # divide combined rax/rdx register by 3
                  # rax = 0xFFFFFFF FFFFFFFF = -5 (quotient)
                  \# rdx = 0xFFFFFFFF FFFFFFFF = -2 (remainder)
```