

# Machine-Level Programming I: Basics

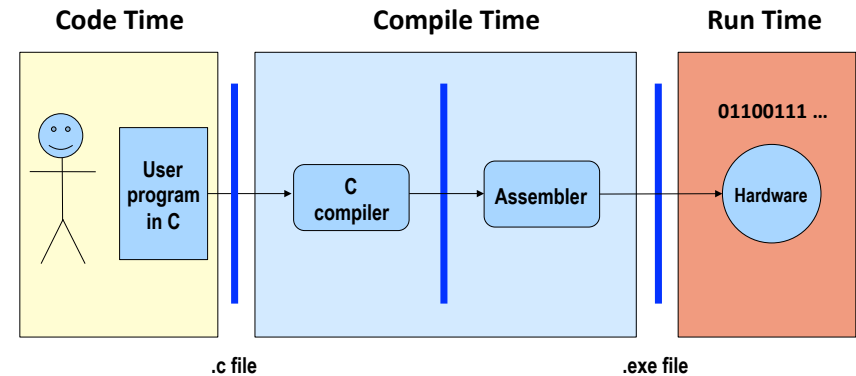
B&O Readings: 3.1-3.5  
CSE 361: Introduction to Systems Software

**Instructor:**  
I-Ting Angelina Lee

Note: these slides were originally created in part by Markus Püschel at Carnegie Mellon University, and in part by Gaetano Borriello and Luis Ceze at University of Washington

1

## Code Translation



What makes programs run fast?

2


## Translation Impacts Performance

- The time required to execute a program depends on:
  - *The program* (as written in C, for instance)
  - *The compiler*: what set of assembler instructions it translates the C program into
  - *The instruction set architecture* (ISA): what set of instructions it makes available to the compiler
  - *The hardware implementation*: how much time it takes to execute an instruction

3

## Today: Machine Programming I: Basics

- What is an ISA (Instruction Set Architecture)
- A really brief history of Intel processors and architectures
- C, assembly, machine code
- Assembly Basics: Registers, operands, move
- Intro to x86-64




4

## Instruction Set Architectures

### ■ The ISA defines:

- The system's state (e.g. registers, memory, program counter)
- The instructions the CPU can execute
- The effect that each of these instructions will have on the system state



5

## General ISA Design Decisions

### ■ Instructions

- What instructions are available? What do they do?
- How are they encoded?

### ■ Registers

- How many registers are there?
- How wide are they?

### ■ Memory

- How do you specify a memory location? (addressing modes)

6

## x86

### ■ Processors that implement the x86 ISA completely dominate the server, desktop and laptop markets

### ■ Evolutionary design

- Backwards compatible up until 8086, introduced in 1978
- Added more features as time goes on

### ■ Complex instruction set computer (CISC)

- Many different instructions with many different formats
  - But, only small subset encountered with Linux programs
- (as opposed to Reduced Instruction Set Computers (RISC), which use simpler instructions)

7

## Intel x86 Evolution: Milestones


<i>Name</i>	<i>Date</i>	<i>Transistors</i>	<i>MHz</i>
■ <b>8086</b>	<b>1978</b>	<b>29K</b>	<b>5-10</b>
■ First 16-bit Intel processor. Basis for IBM PC & DOS			
■ 1MB address space			
■ <b>386</b>	<b>1985</b>	<b>275K</b>	<b>16-33</b>
■ First 32 bit Intel processor , referred to as IA32			
■ Added “flat addressing”, capable of running Unix			
■ <b>Pentium 4E</b>	<b>2004</b>	<b>125M</b>	<b>2800-3800</b>
■ First 64-bit Intel x86 processor, referred to as x86-64			
■ <b>Core 2</b>	<b>2006</b>	<b>291M</b>	<b>1060-3500</b>
■ First multi-core Intel processor			
■ <b>Core i7</b>	<b>2008</b>	<b>731M</b>	<b>1700-3900</b>
■ Four cores			

8

## Intel x86 Processors, cont.

### ■ Machine Evolution

■ 386	1985	0.3M
■ Pentium	1993	3.1M
■ Pentium/MMX	1997	4.5M
■ PentiumPro	1995	6.5M
■ Pentium III	1999	8.2M
■ Pentium 4	2001	42M
■ Core 2 Duo	2006	291M
■ Core i7	2008	731M




### ■ Added Features

- Instructions to support multimedia operations
- Instructions to enable more efficient conditional operations
- Transition from 32 bits to 64 bits
- More cores

9

## Moore's Law



Gordon Moore

*"The number of transistors will double every year", 1965*

*("...or every two years", 1975)*

David House (Intel) says due to transistors' performance improvement, **performance** will double every 18 months.

10

## x86 Clones: Advanced Micro Devices (AMD)

### ■ Historically

- AMD has followed just behind Intel
- A little bit slower, a lot cheaper

### ■ Then

- Recruited top circuit designers from Digital Equipment Corp. and other downward trending companies
- Built Opteron: tough competitor to Pentium 4
- Developed x86-64, their own extension to 64 bits

### ■ Recent Years

- Intel got its act together
  - Leads the world in semiconductor technology
- AMD has fallen behind
  - Relies on external semiconductor manufacturer

11

## Intel's 64-Bit History

### ■ 2001: Intel Attempts Radical Shift from IA32 to IA64

- Totally different architecture (Itanium)
- Executes IA32 code only as legacy
- Performance disappointing

### ■ 2003: AMD Steps in with Evolutionary Solution

- x86-64 (now called "AMD64")

### ■ Intel Felt Obligated to Focus on IA64

- Hard to admit mistake or that AMD is better

### ■ 2004: Intel Announces EM64T extension to IA32

- Extended Memory 64-bit Technology
- Almost identical to x86-64!

### ■ All but low-end x86 processors support x86-64

- But, lots of code still runs in 32-bit mode

12

## Today: Machine Programming I: Basics

- What is an ISA (Instruction Set Architecture)
- History of Intel processors and architectures
- C, assembly, machine code
- Assembly Basics: Registers, operands, move
- Intro to x86-64


13

## Definitions

- **Architecture:** (also ISA: instruction set architecture) The parts of a processor design that one needs to understand or write assembly/machine code.
  - Examples: instruction set specification, registers.
- **Microarchitecture:** Implementation of the architecture.
  - Examples: cache sizes and core frequency.
- **Code Forms:**
  - **Machine Code:** The byte-level programs that a processor executes
  - **Assembly Code:** A text representation of machine code
- **Example ISAs:**
  - Intel: x86, IA32, Itanium, x86-64
  - ARM: Used in almost all mobile phones

14

## Assembly/Machine Code View



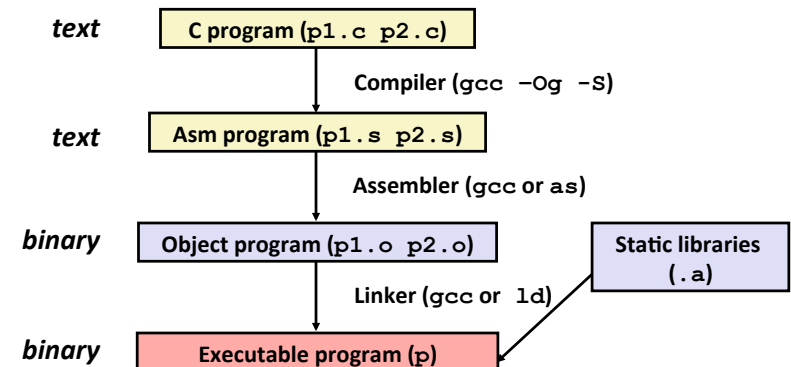
### Programmer-Visible State

- **PC: Program counter**
  - Address of next instruction
  - Called "RIP" (x86-64)
- **Register file**
  - Heavily used program data
- **Condition codes**
  - Store status information about most recent arithmetic or logical operation
  - Used for conditional branching
- **Memory**
  - Byte addressable array
  - Code and user data
  - Stack to support procedures

15

## Turning C into Object Code

- Code in files `p1.c` `p2.c`
- Compile with command: `gcc -Og p1.c p2.c -o p`
  - Use basic optimizations (`-Og`) [New to recent versions of GCC]
  - Put resulting binary in file `p`



16

## Compiling Into Assembly

### C Code (sum.c)

```
long plus(long x, long y);

void sumstore(long x, long y,
              long *dest)
{
    long t = plus(x, y);
    *dest = t;
}
```

### Generated x86-64 Assembly

```
sumstore:
    pushq    %rbx
    movq     %rdx, %rbx
    call     plus
    movq     %rax, (%rbx)
    popq     %rbx
    ret
```

Obtain (on linuxlab machines) with command

```
gcc -Og -S sum.c
```

Produces file `sum.s`

**Warning:** your result may vary due to different compiler versions or platform

17

## Assembly Characteristics: Data Types

- “Integer” data of 1, 2, 4, or 8 bytes
  - Data values
  - Addresses (untyped pointers)
- Floating point data of 4, 8, or 10 bytes
- Code: Byte sequences encoding series of instructions
- No aggregate types such as arrays or structures
  - Just contiguously allocated bytes in memory

18

## Assembly Characteristics: Operations

- Perform arithmetic function on register or memory data
- Transfer data between memory and register
  - Load data from memory into register
  - Store register data into memory
- Transfer control
  - Unconditional jumps to/from procedures
  - Conditional branches

19

## Object Code

### Code for `sumstore`

```
0x0400595:
    0x53
    0x48
    0x89
    0xd3
    0xe8
    0xf2
    0xff
    0xff
    0xff
    0x48
    0x89
    0x03
    0x5b
    0xc3
```

- Total of 14 bytes
- Each instruction 1, 3, or 5 bytes
- Starts at address 0x0400595

### ■ Assembler

- Translates `.s` into `.o`
- Binary encoding of each instruction
- Nearly-complete image of executable code
- Missing linkages between code in different files

### ■ Linker

- Resolves references between files
- Combines with static run-time libraries
  - E.g., code for `malloc`, `printf`
- Some libraries are *dynamically linked*
  - Linking occurs when program begins execution

20

## Machine Instruction Example

```
*dest = t;
```

### ■ C Code

- Store value `t` where designated by `dest`

```
movq %rax, (%rbx)
```

### ■ Assembly

- Move 8-byte value to memory
  - Quad words in x86-64 parlance
- Operands:
  - `t`: Register `%rax`
  - `dest`: Register `%rbx`
  - `*dest`: Memory `M[%rbx]`

```
0x40059e: 48 89 03
```

### ■ Object Code

- 3-byte instruction
- Stored at address `0x40059e`

## Disassembling Object Code

### Disassembled

```
0000000000400595 <sumstore>:
400595: 53          push    %rbx
400596: 48 89 d3    mov     %rdx,%rbx
400599: e8 f2 ff ff callq   400590 <plus>
40059e: 48 89 03    mov     %rax, (%rbx)
4005a1: 5b         pop     %rbx
4005a2: c3         retq
```

### ■ Disassembler

```
objdump -d sum
```

- Useful tool for examining object code
- Analyzes bit pattern of series of instructions
- Produces approximate rendition of assembly code
- Can be run on either a `.out` (complete executable) or `.o` file

21

22

## Alternate Disassembly

### Object

```
0x0400595:
0x53
0x48
0x89
0xd3
0xe8
0xf2
0xff
0xff
0xff
0x48
0x89
0x03
0x5b
0xc3
```

### Disassembled

```
Dump of assembler code for function sumstore:
0x0000000000400595 <+0>: push    %rbx
0x0000000000400596 <+1>: mov     %rdx,%rbx
0x0000000000400599 <+4>: callq   0x400590 <plus>
0x000000000040059e <+9>: mov     %rax, (%rbx)
0x00000000004005a1 <+12>: pop     %rbx
0x00000000004005a2 <+13>: retq
```

### ■ Within gdb Debugger

```
gdb sum
```

```
disassemble sumstore
```

- Disassemble procedure
- `x/14xb sumstore`
  - Examine the 14 bytes starting at `sumstore`

23

## What Can be Disassembled?

```
% objdump -d WINWORD.EXE

WINWORD.EXE: file format pei-i386
```

```
No symbols in "WINWORD.EXE".
Disassembly of section .text:
```

```
30001000 <.text>:
```

```
30001000:
30001001:
30001003:
30001005:
3000100a:
```

**Reverse engineering forbidden by  
Microsoft End User License Agreement**

- Anything that can be interpreted as executable code
- Disassembler examines bytes and reconstructs assembly source

24

## Today: Machine Programming I: Basics

- What is an ISA (Instruction Set Architecture)
- History of Intel processors and architectures
- C, assembly, machine code
- Assembly Basics: Registers, operands, move
- Intro to x86-64

25

## x86-64 Integer Registers

%rax	%eax	%r8	%r8d
%rbx	%ebx	%r9	%r9d
%rcx	%ecx	%r10	%r10d
%rdx	%edx	%r11	%r11d
%rsi	%esi	%r12	%r12d
%rdi	%edi	%r13	%r13d
%rsp	%esp	%r14	%r14d
%rbp	%ebp	%r15	%r15d

- Can reference low-order 4 bytes (also low-order 1 & 2 bytes)

26

## Some History: IA32 Registers

general purpose	%eax	%ax	%ah	%al	accumulate
	%ecx	%cx	%ch	%cl	counter
	%edx	%dx	%dh	%dl	data
	%ebx	%bx	%bh	%bl	base
	%esi	%si			source index
	%edi	%di			destination index
	%esp	%sp			stack pointer
	%ebp	%bp			base pointer

16-bit virtual registers (backwards compatibility)

27

## x86-64 Integer Registers

%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

- Can reference low-order 4 bytes (also low-order 1 & 2 bytes)

28

## Moving Data

### ■ Moving Data

`movq Source, Dest;`  
(Also `movl`, `movw`, `movb`)

### ■ x86-64 can still use 32-bit instructions that generate 32-bit results

- Higher-order bits of destination register are just set to 0
- Example: `addl`

%rax
%rcx
%rdx
%rbx
%rsi
%rdi
%rsp
%rbp
%rN

29

## Moving Data: Operand Types

### ■ **Immediate:** Constant integer data

- Example: `$0x400`, `$-533`
- Like C constant, but prefixed with `'$'`
- Encoded with 1, 2, or 4 bytes

### ■ **Register:** One of 16 integer registers

- Example: `%rax`, `%r13`
- But `%rsp` reserved for special use
- Others have special uses for particular instructions

### ■ **Memory:**

- 8 consecutive bytes of memory at address given by register
  - Have to use the 8-byte form!
- Simplest example: `(%rax)`
- Various other "address modes"

%rax
%rcx
%rdx
%rbx
%rsi
%rdi
%rsp
%rbp
%rN

30

## movq Operand Combinations

	Source	Dest	Src, Dest	C Analog
movq	Imm	Reg	<code>movq \$0x4, %rax</code>	<code>temp = 0x4;</code>
		Mem	<code>movq \$-147, (%rax)</code>	<code>*p = -147;</code>
	Reg	Reg	<code>movq %rax, %rdx</code>	<code>temp2 = temp1;</code>
		Mem	<code>movq %rax, (%rdx)</code>	<code>*p = temp;</code>
	Mem	Reg	<code>movq (%rax), %rdx</code>	<code>temp = *p;</code>

*Cannot do memory-memory transfer with a single instruction*

31

## Simple Memory Addressing Modes

### ■ Normal (R) Mem[Reg[R]]

- Register R specifies memory address
- Aha! Pointer dereferencing in C

`movq (%rcx), %rax`

### ■ Displacement D(R) Mem[Reg[R]+D]

- Register R specifies start of memory region
- Constant displacement D specifies offset

`movq 8(%rbp), %rdx`

32



## Example of Simple Addressing Modes


```
void swap(long *xp,
          long *yp) {
    long t0 = *xp;
    long t1 = *yp;
    *xp = t1;
    *yp = t0;
}
```

```
swap:
    movq    (%rdi), %rax # t0 = *xp
    movq    (%rsi), %rdx # t1 = *yp
    movq    %rdx, (%rdi) # *xp = t1
    movq    %rax, (%rsi) # *yp = t0
    ret
```

33

## Understanding Swap()

```
void swap(long *xp,
          long *yp) {
    long t0 = *xp;
    long t1 = *yp;
    *xp = t1;
    *yp = t0;
}
```



Register	Value
%rdi	xp
%rsi	yp

```
swap:
    movq    (%rdi), %rax # t0 = *xp
    movq    (%rsi), %rdx # t1 = *yp
    movq    %rdx, (%rdi) # *xp = t1
    movq    %rax, (%rsi) # *yp = t0
    ret
```

34

## Understanding Swap()

```
void swap(long *xp,
          long *yp) {
    long t0 = *xp;
    long t1 = *yp;
    *xp = t1;
    *yp = t0;
}
```

Registers		Memory	
%rdi	0x120	123	Address 0x120
%rsi	0x100		0x118
%rax			0x110
%rdx			0x108
		456	0x100

Register	Value
%rdi	xp
%rsi	yp

```
swap:
    movq    (%rdi), %rax # t0 = *xp
    movq    (%rsi), %rdx # t1 = *yp
    movq    %rdx, (%rdi) # *xp = t1
    movq    %rax, (%rsi) # *yp = t0
    ret
```

35

## Understanding Swap()

```
void swap(long *xp,
          long *yp) {
    long t0 = *xp;
    long t1 = *yp;
    *xp = t1;
    *yp = t0;
}
```


Registers		Memory	
%rdi	0x120	123	Address 0x120
%rsi	0x100		0x118
%rax			0x110
%rdx			0x108
		456	0x100

Register	Value
%rdi	xp
%rsi	yp

```
swap:
    movq    (%rdi), %rax # t0 = *xp
    movq    (%rsi), %rdx # t1 = *yp
    movq    %rdx, (%rdi) # *xp = t1
    movq    %rax, (%rsi) # *yp = t0
    ret
```

36


## Understanding Swap()



Register	Value	swap:			
%rdi	xp	movq	(%rdi), %rax	#	t0 = *xp
%rsi	yp	movq	(%rsi), %rdx	#	t1 = *yp
%rax	t0	movq	%rdx, (%rdi)	#	*xp = t1
		movq	%rax, (%rsi)	#	*yp = t0
		ret			

37


## Understanding Swap()



Register	Value	swap:			
%rdi	xp	movq	(%rdi), %rax	#	t0 = *xp
%rsi	yp	movq	(%rsi), %rdx	#	t1 = *yp
%rax	t0	movq	%rdx, (%rdi)	#	*xp = t1
		movq	%rax, (%rsi)	#	*yp = t0
		ret			

38


## Understanding Swap()



Register	Value	swap:			
%rdi	xp	movq	(%rdi), %rax	#	t0 = *xp
%rsi	yp	movq	(%rsi), %rdx	#	t1 = *yp
%rax	t0	movq	%rdx, (%rdi)	#	*xp = t1
%rdx	t1	movq	%rax, (%rsi)	#	*yp = t0
		ret			

39

## Understanding Swap()



Register	Value	swap:			
%rdi	xp	movq	(%rdi), %rax	#	t0 = *xp
%rsi	yp	movq	(%rsi), %rdx	#	t1 = *yp
%rax	t0	movq	%rdx, (%rdi)	#	*xp = t1
%rdx	t1	movq	%rax, (%rsi)	#	*yp = t0
		ret			

40

## Understanding Swap()

```
void swap(long *xp,
          long *yp) {
    long t0 = *xp;
    long t1 = *yp;
    *xp = t1;
    *yp = t0;
}
```

Registers		Memory	
%rdi	0x120	456	Address 0x120
%rsi	0x100		0x118
%rax	123		0x110
%rdx	456	123	0x108
			0x100

Register	Value	swap:	
%rdi	xp	movq	(%rdi), %rax # t0 = *xp
%rsi	yp	movq	(%rsi), %rdx # t1 = *yp
%rax	t0	movq	%rdx, (%rdi) # *xp = t1
%rdx	t1	movq	%rax, (%rsi) # *yp = t0
		ret	

41

## Complete Memory Addressing Modes

### Most General Form

$D(Rb, Ri, S)$        $Mem[Reg[Rb] + S * Reg[Ri] + D]$

- D: Constant "displacement" 1, 2, or 4 bytes
- Rb: Base register: Any of 16 integer registers
- Ri: Index register: Any, except for `%rsp`
- S: Scale: 1, 2, 4, or 8 (*why these numbers?*)

### Special Cases

$(Rb, Ri)$        $Mem[Reg[Rb] + Reg[Ri]]$

$D(Rb, Ri)$        $Mem[Reg[Rb] + Reg[Ri] + D]$

$(Rb, Ri, S)$        $Mem[Reg[Rb] + S * Reg[Ri]]$

42

## Address Computation Examples

%edx	0xf000
%ecx	0x0100

Expression	Address Computation	Address
<code>0x8(%edx)</code>		
<code>(%edx,%ecx)</code>		
<code>(%edx,%ecx,4)</code>		
<code>0x80(,%edx,2)</code>		

## Address Computation Examples

%edx	0xf000
%ecx	0x0100

Expression	Address Computation	Address
<code>0x8(%edx)</code>	<code>0xf000 + 0x8</code>	<code>0xf008</code>
<code>(%edx,%ecx)</code>	<code>0xf000 + 0x100</code>	<code>0xf100</code>
<code>(%edx,%ecx,4)</code>	<code>0xf000 + 4*0x100</code>	<code>0xf400</code>
<code>0x80(,%edx,2)</code>	<code>2*0xf000 + 0x80</code>	<code>0x1e080</code>



43

44

## Today: Machine Programming I: Basics

- History of Intel processors and architectures
- C, assembly, machine code
- Assembly Basics: Registers, operands, move
- Arithmetic & logical operations

45

## Address Computation Instruction

- **leaq** Src, Dst
  - Src is address mode expression (i.e., in the form of D(Rb,Ri,S))
  - Set Dst to address denoted by expression
  - (leaq stands for *load effective address*)
- Uses
  - Computing addresses without a memory reference
    - E.g., translation of `p = &x[i];`
  - Computing arithmetic expressions of the form  $x + k*y$ 
    - $k = 1, 2, 4, \text{ or } 8$
- Example

```
long m12(long x)
{
    return x*12;
}
```

Converted to ASM by compiler:

```
leaq (%rdi,%rdi,2), %rax # t <- x+x*2
salq $2, %rax           # return t<<2
```

46

## Some Arithmetic Operations

- Two Operand Instructions:

Format	Computation	
<b>addq</b>	Src, Dest	Dest = Dest + Src
<b>subq</b>	Src, Dest	Dest = Dest - Src
<b>imulq</b>	Src, Dest	Dest = Dest * Src
<b>salq</b>	Src, Dest	Dest = Dest << Src
<b>sarq</b>	Src, Dest	Dest = Dest >> Src
<b>shrq</b>	Src, Dest	Dest = Dest >> Src
<b>xorq</b>	Src, Dest	Dest = Dest ^ Src
<b>andq</b>	Src, Dest	Dest = Dest & Src
<b>orq</b>	Src, Dest	Dest = Dest   Src

Also called **shlq**  
**Arithmetic**  
**Logical**

- Watch out for argument order!
- No distinction between signed and unsigned int (why?)

47

## Some Arithmetic Operations

- One Operand Instructions

<b>incq</b>	Dest	Dest = Dest + 1
<b>decq</b>	Dest	Dest = Dest - 1
<b>negq</b>	Dest	Dest = - Dest
<b>notq</b>	Dest	Dest = ~Dest

- See book for more instructions:

**movzbw, movzbl, movzwl, movzbq, movzwq**  
**movsbw, movsbl, movswl, movsbq, movswq, movslq**

- Why is there not a **movz1q**?

48

## Arithmetic Expression Example

```
long arith
(long x, long y, long z)
{
    long t1 = x+y;
    long t2 = z+t1;
    long t3 = x+4;
    long t4 = y * 48;
    long t5 = t3 + t4;
    long rval = t2 * t5;
    return rval;
}
```

```
arith:
    leaq    (%rdi,%rsi), %rax
    addq    %rdx, %rax
    leaq    (%rsi,%rsi,2), %rdx
    salq    $4, %rdx
    leaq    4(%rdi,%rdx), %rcx
    imulq    %rcx, %rax
    ret
```

### Interesting Instructions

- **leaq**: address computation
- **salq**: shift
- **imulq**: multiplication
  - But, only used once

49

## Understanding Arithmetic Expression Example

```
long arith
(long x, long y, long z)
{
    long t1 = x+y;
    long t2 = z+t1;
    long t3 = x+4;
    long t4 = y * 48;
    long t5 = t3 + t4;
    long rval = t2 * t5;
    return rval;
}
```

```
arith:
    leaq    (%rdi,%rsi), %rax    # t1
    addq    %rdx, %rax          # t2
    leaq    (%rsi,%rsi,2), %rdx
    salq    $4, %rdx            # t4
    leaq    4(%rdi,%rdx), %rcx   # t5
    imulq    %rcx, %rax          # rval
    ret
```

Register	Use(s)
%rdi	Argument <b>x</b>
%rsi	Argument <b>y</b>
%rdx	Argument <b>z</b>
%rax	t1, t2, rval
%rdx	t4
%rcx	t5

50

## Machine Programming I: Summary

- History of Intel processors and architectures
  - Evolutionary design leads to many quirks and artifacts
- C, assembly, machine code
  - New forms of visible state: program counter, registers, ...
  - Compiler must transform statements, expressions, procedures into low-level instruction sequences
- Assembly Basics: Registers, operands, move
  - The x86-64 move instructions cover wide range of data movement forms
- Arithmetic
  - C compiler will figure out different instruction combinations to carry out computation

51