

The examples and discussion in the following slides have been adapted from a variety of sources, including:

Chapter 3 of Computer Systems 3rd Edition by Bryant and O'Hallaron
x86 Assembly/GAS Syntax on WikiBooks

(http://en.wikibooks.org/wiki/X86_Assembly/GAS_Syntax)

Using Assembly Language in Linux by Phillip ??

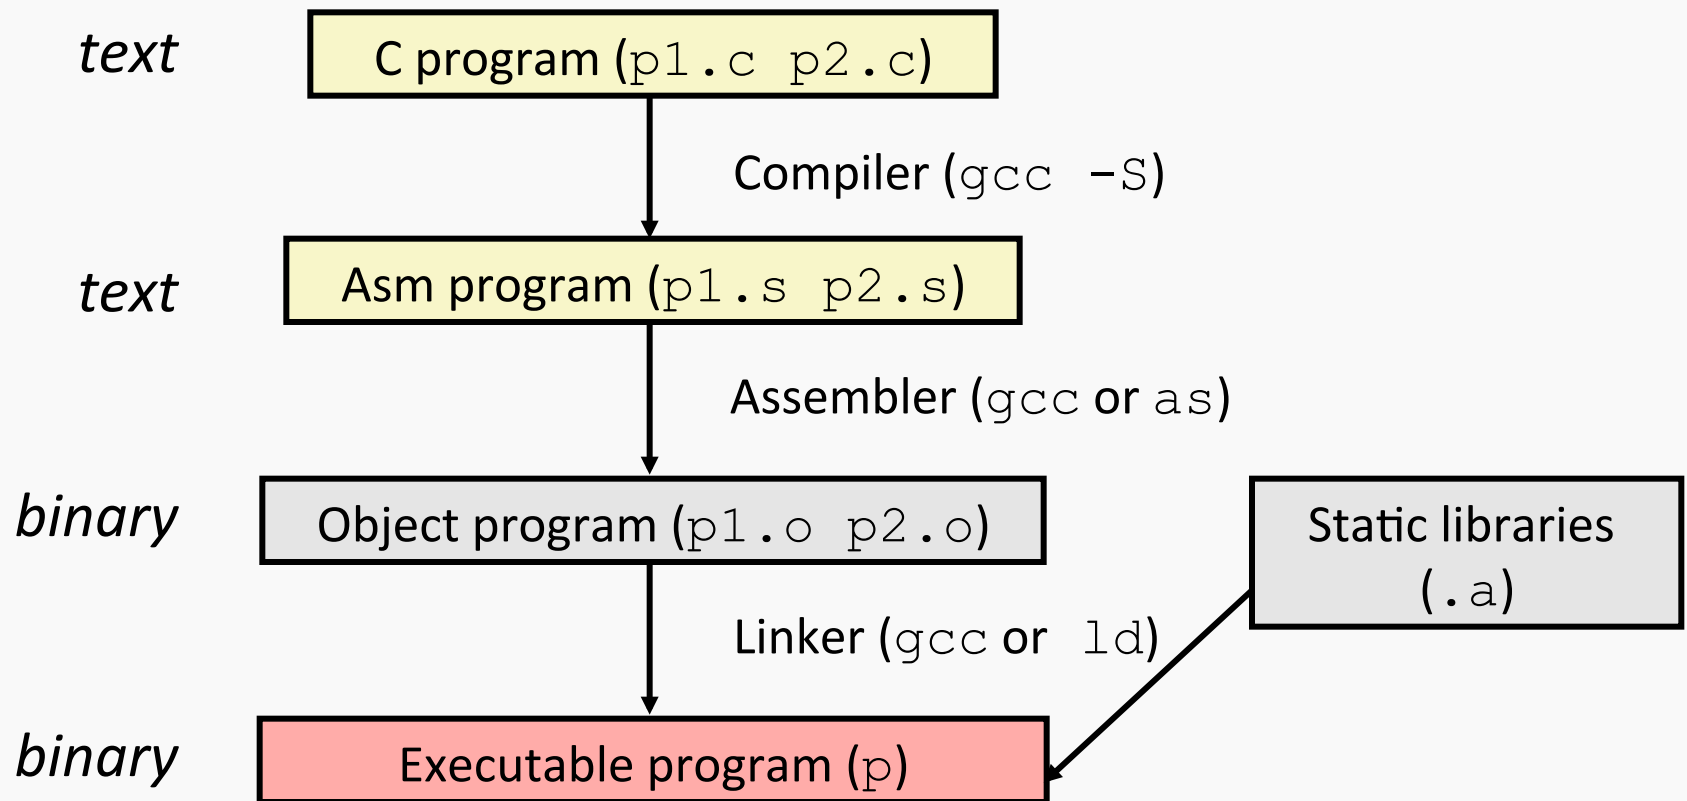
(<http://asm.sourceforge.net/articles/linasm.html>)

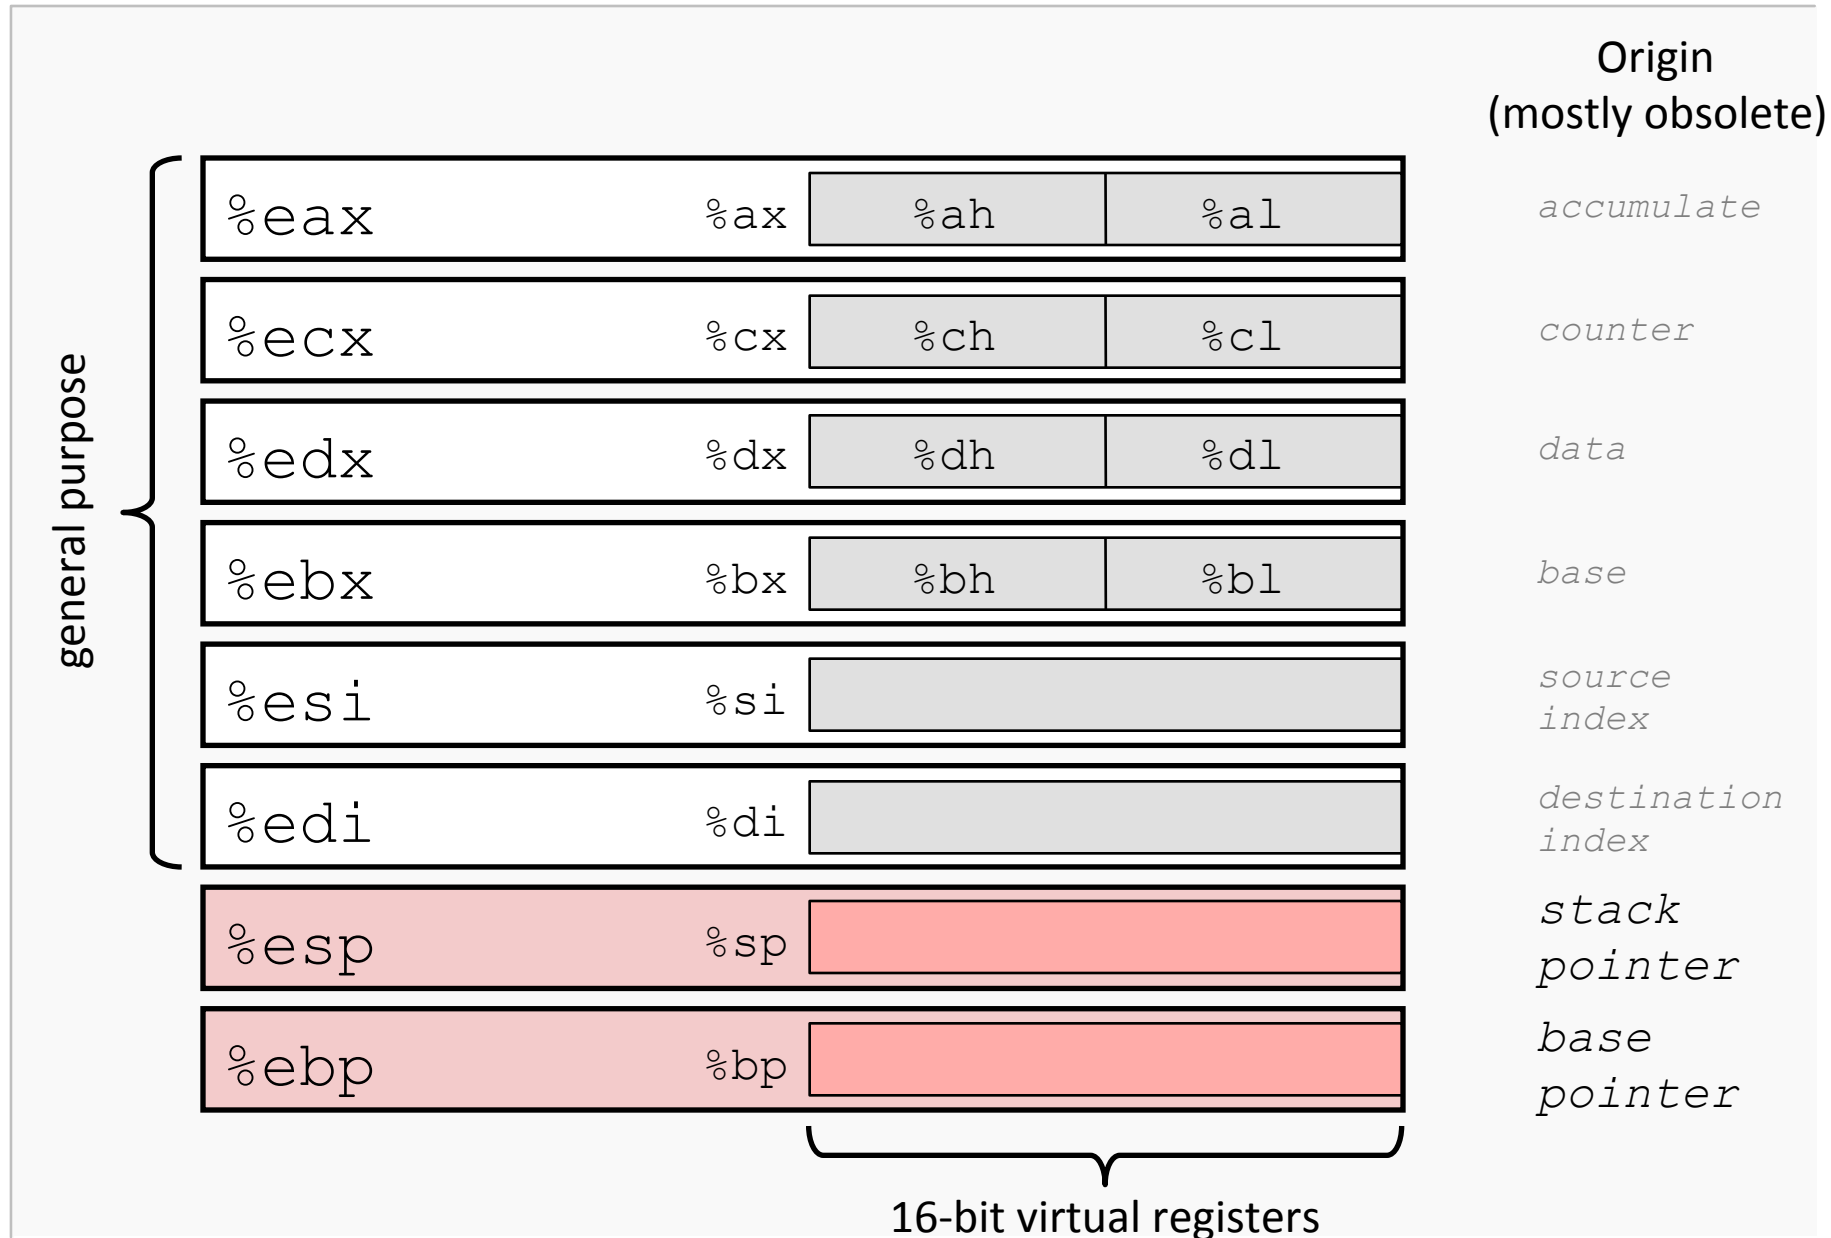
The C code was compiled to assembly with `gcc` version 4.8.3 on CentOS 7.

Unless noted otherwise, the assembly code was generated using the following command line:

```
gcc -S -m64 -fno-asynchronous-unwind-tables -mno-red-zone -O0 file.c
```

AT&T assembly syntax is used, rather than Intel syntax, since that is what the `gcc` tools use.





<code>%rax</code>	<code>%eax</code>
<code>%rbx</code>	<code>%ebx</code>
<code>%rcx</code>	<code>%ecx</code>
<code>%rdx</code>	<code>%edx</code>
<code>%rsi</code>	<code>%esi</code>
<code>%rdi</code>	<code>%edi</code>
<code>%rsp</code>	<code>%esp</code>
<code>%rbp</code>	<code>%ebp</code>

<code>%r8</code>	<code>%r8d</code>
<code>%r9</code>	<code>%r9d</code>
<code>%r10</code>	<code>%r10d</code>
<code>%r11</code>	<code>%r11d</code>
<code>%r12</code>	<code>%r12d</code>
<code>%r13</code>	<code>%r13d</code>
<code>%r14</code>	<code>%r14d</code>
<code>%r15</code>	<code>%r15d</code>

- Extend existing registers. Add 8 new ones.
- Make `%ebp/%rbp` general purpose

Due to the long history of the x86 architecture, the terminology for data lengths can be somewhat confusing:

byte	b	8 bits, no surprises there
short	s	16-bit integer or 32-bit float
word	w	16-bit value
long	l	32-bit integer or 64-bit float (aka double word)
quad	q	64-bit integer

The single-character abbreviations are used in the names of many of the x86 assembly instructions to indicate the length of the operands.

As long as the widths of the operands match, any of these suffixes can be used with the assembly instructions that are discussed in the following slides; for simplicity, we will generally restrict the examples to operations on `long` values.

```
.file "simplest.c"
.text
.globl main
.type main, @function
main:
pushq   %rbp
movq    %rsp, %rbp
subq    $16, %rsp
movl    $5, -4(%rbp)
movl    $16, -8(%rbp)
movl    -8(%rbp), %eax
movl    -4(%rbp), %edx
addl    %edx, %eax
movl    %eax, -12(%rbp)
movl    $0, %eax
popq    %rbp
ret
.size   main, .-main
.ident  "GCC: (GNU) 4.8.3 20140911 (Red Hat 4.8.3-9)"
.section .note.GNU-stack,"",@progbits
```

```
gcc -O0 -S -Wall -m64 simplest.c
```

```
int main() {
    int x, y, t;
    x = 5;
    y = 16;
    t = x + y;
    return 0;
}
```

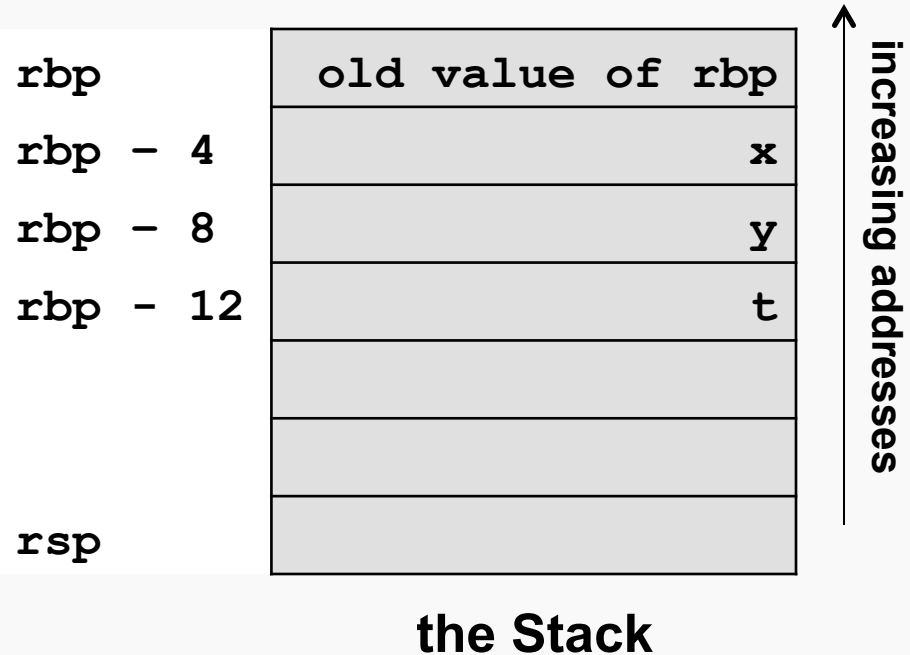
Simple Example: Memory Layout

Local variables and function parameters are stored in memory, and organized in a *stack frame*.

Two registers are used to keep track of the organization:

- rsp address of the top element on the stack
- rbp address of the first element in the current stack frame

```
int main() {  
  
    int x, y, t;  
  
    x = 5;  
    y = 16;  
    t = x + y;  
  
    return 0;  
}
```



Many machine-level operations require that data be transferred between memory and registers.

The most basic instructions for this are the variants of the `mov` instruction:

```
movl src, dest
      dest := src
```

This copies a 32-bit value from *src* into *dest*. `movq` moves 64 bit values in the same fashion.

Despite the name, it has no effect on the value of *src*.

The two operands can be specified in a number of ways:

- immediate values
- one of the 16 x86-64 integer registers (or their virtual registers)
- memory address

Immediate: Constant integer data

Example: `$0x400`, `$-533`

Like C constant, but prefixed with `'$'`

Encoded with 1, 2, or 4 bytes

Register: One of the 16 integer registers

Example: `%eax`, `%edx` (reg names preceded by `'%'`)

But `%rsp` and `%rbp` reserved for special use

Others have special uses for particular instructions

Memory: N consecutive bytes of memory at address given by register, N is specified by the instruction name, `movl` = 4 bytes, `movq` = 8 bytes.

Simplest example: `(%rax)`

Various other “address modes”

X86-64 assembly

```
movl $0x10, %eax
```

```
movl $42, %ebx
```

```
movl %ecx, %edx
```

```
movl %eax, (%rbx)
```

```
movl (%rbx), %eax
```

```
movl -4(%rbp), %eax
```

C analog

```
a = 16;
```

```
b = 42;
```

```
d = c;
```

```
*b = a
```

```
a = *b
```

```
a = *(rbp - 4)
```

Mapping:

	reg
a	%eax
b	%ebx
c	%ecx
d	%edx

```
int main() {  
    int x, y, t;  
    x = 5;  
    y = 16;  
    t = x + y;  
    return 0;  
}
```

Registers

eax	
ebx	
ecx	
edx	
edi	
esi	

rbp
rbp - 4
rbp - 8
rbp - 12

old value of rbp	
	x
	y
	t

the Stack

```
movl $5, -4(%rbp)
```

```
movl $16, -8(%rbp)
```

```
movl -8(%rbp), %eax  
movl -4(%rbp), %edx  
addl %edx, %eax  
movl %eax, -12(%rbp)
```

```
int main() {  
    int x, y, t;  
    x = 5;  
    y = 16;  
    t = x + y;  
    return 0;  
}
```

rbp
rbp - 4
rbp - 8
rbp - 12

old value of rbp	
	5
	??
	??

the Stack

```
movl $5, -4(%rbp)
```

Registers

eax	??
edx	??

```
int main() {  
  
    int x, y, t;  
    x = 5;  
    y = 16;  
    t = x + y;  
    return 0;  
}
```

eax	??
edx	??

rbp	old value of rbp
rbp - 4	5
rbp - 8	16
rbp - 12	??

the Stack

```
movl $5, -4(%rbp)  
movl $16, -8(%rbp)
```

```
int main() {  
  
    int x, y, t;  
    x = 5;  
    y = 16;  
    t = x + y;  
    return 0;  
}
```

rbp	old value of rbp
rbp - 4	5
rbp - 8	16
rbp - 12	??

the Stack

Registers	
eax	16
edx	??

```
movl $5, -4(%rbp)
```

```
movl $16, -8(%rbp)
```

```
movl -8(%rbp), %eax  
movl -4(%rbp), %edx  
addl %edx, %eax  
movl %eax, -12(%rbp)
```

```
int main() {
    int x, y, t;
    x = 5;
    y = 16;
    t = x + y;
    return 0;
}
```

Registers

eax	16
edx	5

rbp
rbp - 4
rbp - 8
rbp - 12

old value of rbp	
	5
	16
	??

the Stack

```
movl $5, -4(%rbp)
```

```
movl $16, -8(%rbp)
```

```
movl -8(%rbp), %eax
movl -4(%rbp), %edx
addl %edx, %eax
movl %eax, -12(%rbp)
```

```
int main() {  
  
    int x, y, t;  
    x = 5;  
    y = 16;  
    t = x + y;  
    return 0;  
}
```

Registers	
eax	21
edx	5

	old value of rbp
rbp	
rbp - 4	5
rbp - 8	16
rbp - 12	??

the Stack

```
movl $5, -4(%rbp)
```

```
movl $16, -8(%rbp)
```

```
movl -8(%rbp), %eax  
movl -4(%rbp), %edx  
addl %edx, %eax  
movl %eax, -12(%rbp)
```


We have the expected addition operation:

```
addl rightop, leftop  
      leftop = leftop + rightop
```

The operand ordering shown here is probably confusing:

- As usual, the destination is listed second.
- But, that's also the first (left-hand) operand when the arithmetic is performed.

This same pattern is followed for all the binary integer arithmetic instructions.

See the discussion of AT&T vs Intel syntax later in the notes for an historical perspective on this.

```
int main() {
    int x, y, t;
    x = 5;
    y = 16;
    t = x + y;
    return 0;
}
```

Registers

eax	21
edx	5

rbp	old value of rbp
rbp - 4	5
rbp - 8	16
rbp - 12	21

the Stack

```
movl $5, -4(%rbp)
movl $16, -8(%rbp)
movl -8(%rbp), %eax
movl -4(%rbp), %edx
addl %edx, %eax
movl %eax, -12(%rbp)
```

In addition:

```
subl rightop, leftop  
    leftop = leftop - rightop
```

```
imull rightop, leftop  
    leftop = leftop * rightop
```

```
negl op  
    op = -op
```

```
incl op  
    op = op + 1
```

```
decl op  
    op = op - 1
```

(Yes, there is a division instruction, but its interface is confusing and we will not need it.)