

CSE 351 Section 4 – GDB and x86-64 Assembly

Hi there! Welcome back to section, we're happy that you're here ☺

The GNU Debugger (GDB)

The GNU Debugger is a powerful debugging tool that will be critical to Lab 2 and Lab 3 and is a useful tool to know as a programmer moving forward. There are tutorials and reference sheets available on the course webpage, but the following tutorial should get you started with the basics:

GDB Tutorial:

- 1) Download `calculator.c` from the class webpage if you didn't already have it from Section 1:
> `wget https://courses.cs.washington.edu/courses/cse351/17au/sections/01/code/calculator.c`
- 2) Compile the file *with debugging symbols* (`-g` flag):
> `gcc -g -o calculator calculator.c`
- 3) Load the binary (executable) into GDB. This will spit out a bunch of information (e.g. version, license).
> `gdb calculator`
- 4) Inside of GDB, use the run command (**run** or just **r**) to execute your program. By default, this will continue until an error or breakpoint is encountered or your program exits.
 - a. Command-line arguments can be passed as additional arguments to **run**:
(gdb) `run 3 4 +`
 - b. To step through the program starting at `main()` instead, use the start command (**start** or just **sta**):
(gdb) `start`

- 5) To view *source* code while debugging, use the list command (**list** or just **l**).
 - a. You can give list a function name ("list <function>") to look at the beginning of a function.
(gdb) `list main`
 - b. You can give list a line number ("list <line>") to look at the lines *around* that line number, or provide a specific range ("list <start>, <end>").
(gdb) `list 45`
(gdb) `list 10, 15`
 - c. "**list**" will display the next 10 lines of code *after* whatever was last displayed and "**list -**" will display the previous 10 lines of code before whatever was last displayed. C

- 6) To view *assembly* code while debugging, use the disassemble command (**disassemble** or just **disas**).
 - a. "**disas**" will display the disassembly of the current function that you are in.
 - b. You can also disassemble specific functions.
(gdb) `disas main`
(gdb) `disas print_operation`

x86-64

- 7) Create breakpoints using the break command (**break** or **b**)
 - a. A breakpoint will stop program execution *before* the shown instruction has been executed!
 - b. You can create a breakpoint at a function name, source code line number, or assembly instruction address. The following all break at the same place:
(gdb) `break main`
(gdb) `break 34`
(gdb) `break *0x4005d5`

- c. Each break point has an associated number. You can view your breakpoints using the info command (**info** or just **i**) and then enable (**enable** or just **en**) or disable (**disable** or just **dis**) specific ones.

```
(gdb) info break
(gdb) disable 3
(gdb) enable 3
```

- 8) Navigating source code within GDB is done while program execution is started (**run** or **start**), but halted (e.g. at a breakpoint).

- a. Use the next command (**next** or just **n**) to execute the next # of lines of *source* code and then break again. This will complete (“step *over*”) any function calls found in the lines of code.

```
(gdb) next
(gdb) next 4
```

- b. Use the step command (**step** or just **s**) to execute the next # of lines of *source* code and then break again. This will step *into* any function calls found in the lines of code.

```
(gdb) step
(gdb) step 4
```

C

- c. Use the “next instruction” command (**nexti** or just **ni**) to execute the next # of lines of *assembly* code and then break again. This will complete (“step *over*”) any function calls.

```
(gdb) nexti
(gdb) nexti 4
```

- d. Use the “step instruction” command (**stepi** or just **si**) to execute the next # of lines of *assembly* code and then break again. This will step *into* any function calls.

```
(gdb) stepi
(gdb) stepi 4
```

x86-64

- e. Use the finish command (**finish** or just **fin**) to step *out* of the current function call.

- f. Use the continue command (**continue** or just **c**) to resume continuous program execution (until next breakpoint is reached or your program terminates).

- 9) You can print the current value of variables or expressions using the print command (**print** or just **p**):

- a. The print command can take an optional format specifier: **/x** (hex), **/d** (decimal), **/u** (unsigned), **/t** (binary), **/c** (char), **/f** (float)

```
(gdb) print /t argc
(gdb) print /x argv
(gdb) print /d argc*2+5
(gdb) print /x $rax
```

- b. The display command (**display** or just **disp**) is similar, but causes the expression to print in the specified format *every time* the program stops.

- 10) You can terminate the current program run using the kill command (**kill** or just **k**). This will allow you to restart execution (run or start) with your breakpoints intact.

- 11) You can exit GDB by either typing **Ctrl-D** or using the quit command (**quit** or just **q**)

x86-64 Assembly Language

Assembly language is a human-readable representation of machine code instructions (generally a one-to-one correspondence). Assembly is machine-specific because the computer architecture and hardware are designed to execute a particular machine code instruction set.

x86-64 is the primary 64-bit instruction set architecture (ISA) used by modern personal computers. It was developed by Intel and AMD and its 32-bit predecessor is called IA32. x86-64 is designed for complex instruction set computing (CISC), generally meaning it contains a larger set of more versatile and more complex instructions.

For this course, we will utilize only a small subset of x86-64's instruction set and omit floating point instructions.

x86-64 Instructions

The subset of x86-64 instructions that we will use in this course take either one or two operands, usually in the form: `operation operand1, operand2`. There are three options for operands:

- **Immediate**: constant integer data (*e.g.* `$0x400`, `$-533`) or an address/label (*e.g.* `Loop`, `main`)
- **Register**: use the data stored in one of the 16 general purpose registers or subsets (*e.g.* `%rax`, `%edi`)
- **Memory**: use the data at the memory address specified by the addressing mode `D(Rb, Ri, S)`

The operation determines the effect of the operands on the processor state and has a suffix ("b" for byte, "w" for word, "l" for long, "q" for quad word) that determines the bit width of the operation. Sometimes the operation size can be inferred from the operands, so the suffix is omitted for brevity.

Control Flow and Condition Codes

Internally, condition codes (Carry, Zero, Sign, Overflow) are set based on the result of the previous operation. The `je*` and `set*` families of instructions use the values of these "flags" to determine their effects. See the table provided on your reference sheet for equivalent conditionals.

An *indirect jump* is specified by adding an asterisk (*) in front of a memory operand and causes your program counter to load the address stored at the computed address.

Procedure Basics

The instructions `push`, `pop`, `call`, and `ret` move the stack pointer (`%rsp`) automatically.

`%rax` is used for the return value and the first six arguments go in `%rdi`, `%rsi`, `%rdx`, `%rcx`, `%r8`, `%r9`
(**"Diane's Silk Dress Cost \$89"**).

Exercises:

1. [CSE351 Au14 Midterm] Symbolically, what does the following code return?

```
movl  (%rdi), %eax          # %rdi -> x
leal  (%eax,%eax,2), %eax   # %rax -> r
addl  %eax, %eax
andl  %esi, %eax           # %rsi -> y
subl  %esi, %eax
ret
```

2. [CSE351 Au15 Midterm] Convert the following C function into x86-64 assembly code. You are not being judged on the efficiency of your code – just the correctness.

```
long happy(long *x, long y, long z) {
    if (y > z)
        return z + y;
    else
        return *x;
}
```

3. Write an equivalent C function for the following x86-64 code:

```
mystery:
    testl    %edx, %edx
    js      .L3
    cmpl    %esi, %edx
    jge    .L3
    movslq  %edx, %rdx
    movl    (%rdi,%rdx,4), %eax
    ret
.L3:
    movl    $0, %eax
    ret
```

4. [Adapted from CSE351 Wi16 Midterm] Suppose before the assembly below is executed, the value of `%rsp` is `0xFFFF8888`.

```
0x00002f:  pushq $7
0x000031:  pushq $5
0x000033:  addq $2, 8(%rsp)
0x000039:  callq someOtherFunction
0x00003e:  ...
```

Immediately *after* the `callq` instruction executes:

a. What is the value of `%rsp` in hexadecimal?

b. Fill in the contents of the stack from `%rsp` (your answer to part a) up to (but not including) `0xFFFF8888`. Fill in the boxes below using hexadecimal. You may not need all rows.

Address	Data
0xFFFF8888	<unknown>

Au16 Midterm Q5

Question 5: The Stack [12 pts]

The recursive factorial function `fact()` and its x86-64 disassembly is shown below:

```
int fact(int n) {
    if(n==0 || n==1)
        return 1;
    return n*fact(n-1);
}
```

```
000000000040052d <fact>:
40052d: 83 ff 00      cmpl  $0, %edi
400530: 74 05        je    400537 <fact+0xa>
400532: 83 ff 01      cmpl  $1, %edi
400535: 75 07        jne  40053e <fact+0x11>
400537: b8 01 00 00 00 movl  $1, %eax
40053c: eb 0d        jmp   40054b <fact+0x1e>
40053e: 57          pushq %rdi
40053f: 83 ef 01      subl  $1, %edi
400542: e8 e6 ff ff ff call  40052d <fact>
400547: 5f          popq  %rdi
400548: 0f af c7      imull %edi, %eax
40054b: f3 c3        rep ret
```

(A) Circle one: [1 pt] `fact()` is saving `%rdi` to the Stack as a **Caller** // **Callee**

(B) How much space (in bytes) does this function take up in our final executable? [2 pt]

(C) **Stack overflow** is when the stack exceeds its limits (i.e. runs into the Heap). Provide an argument to `fact(n)` here that will cause stack overflow. [2 pt]

«Problem continued on next page»

- (D) If we use the main function shown below, answer the following for the execution of the entire program: [4 pt]

```
void main() {  
    printf("result = %d\n", fact(3));  
}
```

Total frames created:	Maximum stack frame depth:
--------------------------	-------------------------------

- (E) In the situation described above where `main()` calls `fact(3)`, we find that the word `0x2` is stored on the Stack at address `0x7fffdc7ba888`. At what address on the Stack can we find the return address to `main()`? [3 pt]

CSE 351 Reference Sheet (Midterm)

Binary	Decimal	Hex
0000	0	0
0001	1	1
0010	2	2
0011	3	3
0100	4	4
0101	5	5
0110	6	6
0111	7	7
1000	8	8
1001	9	9
1010	10	A
1011	11	B
1100	12	C
1101	13	D
1110	14	E
1111	15	F

2^0	2^1	2^2	2^3	2^4	2^5	2^6	2^7	2^8	2^9	2^{10}
1	2	4	8	16	32	64	128	256	512	1024

IEEE 754 FLOATING-POINT STANDARD


Value: $\pm 1 \times \text{Mantissa} \times 2^{\text{Exponent}}$

Bit fields: $(-1)^S \times 1.M \times 2^{(E-\text{bias})}$

where Single Precision Bias = 127,
Double Precision Bias = 1023.

IEEE Single Precision and

Double Precision Formats:



IEEE 754 Symbols

Exponent	Fraction	Object
0	0	± 0
0	$\neq 0$	\pm Denorm
1 to MAX - 1	anything	\pm Fl. Pt. Num.
MAX	0	$\pm\infty$
MAX	$\neq 0$	NaN

S.P. MAX = 255, D.P. MAX = 2047

Assembly Instructions

<code>mov a, b</code>	Copy from a to b.
<code>movs a, b</code>	Copy from a to b with sign extension. Needs <i>two</i> width specifiers.
<code>movz a, b</code>	Copy from a to b with zero extension. Needs <i>two</i> width specifiers.
<code>lea a, b</code>	Compute address and store in b. <i>Note:</i> the scaling parameter of memory operands can only be 1, 2, 4, or 8.
<code>push src</code>	Push <code>src</code> onto the stack and decrement stack pointer.
<code>pop dst</code>	Pop from the stack into <code>dst</code> and increment stack pointer.
<code>call <func></code>	Push return address onto stack and jump to a procedure.
<code>ret</code>	Pop return address and jump there.
<code>add a, b</code>	Add from a to b and store in b (and sets flags).
<code>sub a, b</code>	Subtract a from b (compute $b-a$) and store in b (and sets flags).
<code>imul a, b</code>	Multiply a and b and store in b (and sets flags).
<code>and a, b</code>	Bitwise AND of a and b, store in b (and sets flags).
<code>sar a, b</code>	Shift value of b <i>right (arithmetic)</i> by a bits, store in b (and sets flags).
<code>shr a, b</code>	Shift value of b <i>right (logical)</i> by a bits, store in b (and sets flags).
<code>shl a, b</code>	Shift value of b <i>left</i> by a bits, store in b (and sets flags).
<code>cmp a, b</code>	Compare b with a (compute $b-a$ and set condition codes based on result).
<code>test a, b</code>	Bitwise AND of a and b and set condition codes based on result.
<code>jmp <label></code>	Unconditional jump to address.
<code>j* <label></code>	Conditional jump based on condition codes (<i>more on next page</i>).
<code>set* a</code>	Set byte based on condition codes.

Conditionals

Instruction		(op) s, d	test a, b	cmp a, b
je	“Equal”	d (op) s == 0	b & a == 0	b == a
jne	“Not equal”	d (op) s != 0	b & a != 0	b != a
js	“Sign” (negative)	d (op) s < 0	b & a < 0	b-a < 0
jns	(non-negative)	d (op) s >= 0	b & a >= 0	b-a >= 0
jg	“Greater”	d (op) s > 0	b & a > 0	b > a
jge	“Greater or equal”	d (op) s >= 0	b & a >= 0	b >= a
jl	“Less”	d (op) s < 0	b & a < 0	b < a
jle	“Less or equal”	d (op) s <= 0	b & a <= 0	b <= a
ja	“Above” (unsigned >)	d (op) s > 0U	b & a < 0U	b > a
jb	“Below” (unsigned >)	d (op) s < 0U	b & a > 0U	b < a

Registers

Name	Convention	Name of “virtual” register		
		Lowest 4 bytes	Lowest 2 bytes	Lowest byte
%rax	Return value – Caller saved	%eax	%ax	%al
%rbx	Callee saved	%ebx	%bx	%bl
%rcx	Argument #4 – Caller saved	%ecx	%cx	%cl
%rdx	Argument #3 – Caller saved	%edx	%dx	%dl
%rsi	Argument #2 – Caller saved	%esi	%si	%sil
%rdi	Argument #1 – Caller saved	%edi	%di	%dil
%rsp	Stack Pointer	%esp	%sp	%spl
%rbp	Callee saved	%ebp	%bp	%bpl
%r8	Argument #5 – Caller saved	%r8d	%r8w	%r8b
%r9	Argument #6 – Caller saved	%r9d	%r9w	%r9b
%r10	Caller saved	%r10d	%r10w	%r10b
%r11	Caller saved	%r11d	%r11w	%r11b
%r12	Callee saved	%r12d	%r12w	%r12b
%r13	Callee saved	%r13d	%r13w	%r13b
%r14	Callee saved	%r14d	%r14w	%r14b
%r15	Callee saved	%r15d	%r15w	%r15b

Sizes

C type	x86-64 suffix	Size (bytes)
char	b	1
short	w	2
int	l	4
long	q	8