

Chapter 4 Mixing Assembly and C¹

The goal is to learn how to use C and assembly in the same program and become aware of performance issues. We will implement a 64-bit multiplication. The related files are in the following directory:

`/home/TDDI11/lab/skel/lab2`

Please copy them as usual to your own account. Your modifications go to “`llmultiply.asm`” and “`main.c`”. The rest of the files are the usual files that we used before.

1.1 Assignment 1, Assembly Implementation

Processors used in a small (embedded) part in a mass produced product have to be very cheap. Therefore, they are quite rudimentary. For example, they seldom have floating point arithmetic capability, or they may be limited to 8-bit integer operations. Software must compensate for the limited hardware capability. We will look at an example in which we want to multiply two numbers larger than the hardware supports natively. On an 8-bit processor this could involve multiplication of two 16-bit integers. As our target already can do that we will look at multiplying two 64-bit operands on our target machine that natively supports only 32-bit operands.

Our compiler supports 64-bit integers with a data type called “long long int”. Since the registers of the Intel processor are only 32-bits wide, how does the compiler generate code to implement $a \times b$ when a and b are long long int’s? This is discussed later on in this chapter.

Write a function in assembly that has the following C signature:

```
void llmultiply(unsigned long long int a, unsigned long long int b, unsigned char *result);
```

The function multiplies the two 64-bit parameters a and b . The result of the multiplication is a 128-bit number. It has to be copied to the array of 16 bytes that is pointed to by `result`. Note that the x86 machines are little-endian, meaning that the least significant byte of a multi-byte number is placed at the lower address, and the most significant byte of a multi-byte number is placed at the higher address.

Based on the explanations in section 4.4, you will implement your multiplication function and test it with tests given in section 4.5. These tests are also given in “`main.c`”. Test your function with at least the given test cases. The given test cases are tailored to generate carries in all possible steps of the calculation.

¹ Note that the structure of the chapters are not identical. Please read them from beginning to the end, before you get to work. Some concepts may be explained at the end of a chapter.

1.2 Assignment 2, C Implementation

Implement the same function, but this time in C. Here you must make sure to use appropriate data type for the multiplication and addition in order to be able to store the entire result in a sufficiently large type.

- Compile it **without** any optimization and verify that it gives correct results.
- Compile it **with** optimization turned on and verify that it gives correct results.

The compiler optimization options are described in section 4.7.

1.3 Assignment 3, Performance Comparison

Call “llmultiply” from a C program in which you test the function. Put the invocation of the function in a loop in order to invoke it many times. Obtain the time or the contents of the CPU cycle register before entering and after exiting the loop. Print the difference of the values before and after the loop. When emulating on Qemu, depending on the host machine, the results might not make sense. The important thing is that you write the program correctly. No worries if Qemu returns wrong values.

Test both your C-version (optimized and non-optimized) and your Assembly-version in the same way (the same number of iterations). Which version is most effective? How big improvement does compiler optimization give?

Note that depending on the computer-architecture and the compiler, the optimization benefits might be partly resulted from optimized loop implementation that affects the testing-loop execution not necessarily the actual instruction inside the loop.

1.4 Multiplication theory

Consider that each 64-bit operand can be split in two 32-bit parts, one contains the high order bits and the other one contains the low order bits.

$$a = a_h \times 2^{32} + a_l$$

$$b = b_h \times 2^{32} + b_l$$

If we break the operands to two parts as suggested above and perform the multiplication with these separated parts we get:

$$\begin{aligned} a \times b &= (a_h \times 2^{32} + a_l) \times (b_h \times 2^{32} + b_l) = \\ &= a_h \times b_h \times 2^{64} + \\ &\quad + (a_h \times b_l + a_l \times b_h) \times 2^{32} + \\ &\quad + a_l \times b_l \end{aligned}$$

All multiplications that appear in the above expressions are now 32-bit multiplications and therefore they can be implemented on an Intel x86 processor. Your only concern is to handle the additions and carries that may appear after the additions. Note that each 32-bit multiplication yields a 64-bit result, but you can only add 32-bits in each addition. A graphical view of the procedure is given below. Think carefully of

where carries can occur. In the figure below, in the lower box that shows the results, the numbers represents a numeration of the bytes. For example “result 3 .. 0” represents 4 bytes that have the lowest significant bits.

$(a_h * b_h)_h$		$(a_h * b_h)_l$	
		$(a_h * b_l)_h$	$(a_h * b_l)_l$
		$(a_l * b_h)_h$	$(a_l * b_h)_l$
+		$(a_l * b_l)_h$	$(a_l * b_l)_l$
result 15.. 12		result 11.. 8	result 7 .. 4
		result 3 .. 0	

Be careful of the carries when performing the additions.

1.5 Test Cases

We provide the following test cases. You should of course add your own. All digits are hexadecimal.

```
0000111122223333 * 0000555566667777 = 0000000005B061D958BF0ECA7C0481B5
3456FEDCAAAA1000 * EDBA00112233FF01 = 309A912AF7188C57E62072DD409A1000
FFFFFFFFDDDDCCCC * BBBBAAA99998888 = BBBB9E2692C5DDCC28F7531048D2C60
FFFFFFFFFFFFFFFF * FFFFFFFFFFFFFFFF = FFFFFFFFFFFFFFFF0000000000000001
00000001FFFFFFFF * 00000001FFFFFFFF = 0000000000000003FFFFFFFFC00000001
FFFFFFFFFFFFFFFF * FFFF0001FFFFFFFF = FFFE0002FFDFFFE0001FFFE00000001
```

1.6 C Function Call Interface

When writing the assembly code that cooperates with C-code you have to follow the compilers idea of how to pass parameters to a function. The parameters are passed on the stack. The stack grows from large addresses to small ones. The last parameter to a function is pushed first on the stack. Hence, the last parameter will be at a larger address than the first parameter. The stack pointer points to the top of the stack and not to the first free location! That means that pushing a value on the stack first decrements the stack pointer and then writes the value.

To understand better, look at the following snapshot of a stack frame just after entering the function:

```
byte 0 of return address | 0x3fffffe8 <-- stack top (esp)
byte 1 of return address | 0x3fffffe9
byte 2 of return address | 0x3fffffea
byte 3 of return address | 0x3fffffeb
;; The first parameter (a) start here.
;; Notice the 32-bit little endianness:
;; The least significant byte come first (byte 0)
;; The most significant byte come last (byte 3)
byte 0 of a, byte 0 of a1 | 0x3fffffec
byte 1 of a, byte 1 of a1 | 0x3fffffed
byte 2 of a, byte 2 of a1 | 0x3fffffee
byte 3 of a, byte 3 of a1 | 0x3fffffef
```

```

byte 4 of a, byte 0 of ah      | 0x3fffffff0
byte 5 of a, byte 1 of ah      | 0x3fffffff1
byte 6 of a, byte 2 of ah      | 0x3fffffff2
byte 7 of a, byte 3 of ah      | 0x3fffffff3
;; The second parameter (b) start here.
byte 0 of b, byte 0 of bl      | 0x3fffffff4
byte 1 of b, byte 1 of bl      | 0x3fffffff5
byte 2 of b, byte 2 of bl      | 0x3fffffff6
byte 3 of b, byte 3 of bl      | 0x3fffffff7
byte 4 of b, byte 0 of bh      | 0x3fffffff8
byte 5 of b, byte 1 of bh      | 0x3fffffff9
byte 6 of b, byte 2 of bh      | 0x3fffffffa
byte 7 of b, byte 3 of bh      | 0x3fffffffb
;; The third parameter (c) start here.
;; Notice that only the address to the array is passed.
byte 0 of result array address| 0x3fffffffc
byte 1 of result array address| 0x3fffffffd
byte 2 of result array address| 0x3fffffffe
byte 3 of result array address| 0x3fffffff -- stack bottom

```

Typically a function has the following prologue:

```

push ebp      ;; save the value of ebp register on the stack
mov ebp, esp  ;; save the address of the stack frame
              ;; (the value of esp after entering the function)

```

The reason for the prologue is to get a fix base pointer for convenient access to the function parameters (the stack pointer esp may be changed in the function to store local variables). It also makes printing of a stack trace easy, which is an important debug feature. After the prologue ebp will contain the value 0x3ffffffe4, i.e. the value of the stack pointer before pushing ebp, that is 0x3ffffffe8, minus the four locations occupied by ebp. The stack will now look like this:

```

byte 0 of previous stack frame (ebp) | 0x3ffffffe4 <-- top (ebp, esp)
byte 1 of previous stack frame (ebp) | 0x3ffffffe5
byte 2 of previous stack frame (ebp) | 0x3ffffffe6
byte 3 of previous stack frame (ebp) | 0x3ffffffe7
byte 0 of return address              | 0x3ffffffe8 <-- previous top
byte 1 of return address              | 0x3ffffffe9
byte 2 of return address              | 0x3ffffffea
byte 3 of return address              | 0x3ffffffeb
byte 0 of a, byte 0 of al             | 0x3ffffffec <-- parameter 1
...

```

To illustrate how you can get that said convenient access to the parameters, let's add the values of b_h and a_l as an example (this addition is of course irrelevant for the assignment). You will find b_h 20 bytes from the address in ebp (count in the picture). To get the value of b_h to register eax we write:

```
mov eax, [ebp + 20]
```

Then we fetch a_l , but as eax now is occupied we have to use another destination register:

```
mov ebx, [ebp + 8]
```

And to add the two (with the result in `eax`) we do:

```
add eax, ebx
```

To make your code more readable it is good to use symbolic names for the offsets:

```
mov eax, [ebp + BH_OFFSET]
mov ebx, [ebp + AL_OFFSET]
add eax, ebx
```

As an exercise, think of how to load the address of the first byte of the result array in `ebx`.

Since we push `ebp` in the prologue inside the function we should restore the previous value before we return. We also have to make sure the stack pointer point to the same address it had when we entered the function, in order to use the correct return address. This is the function epilogue that is to pop `ebp` from the stack before leaving the function.

```
mov esp, ebp ;; perhaps needed...
pop ebp
ret
```

1.7 Compiler Optimizations

In order to compile without any optimization, pass the `-O0` switch (minus capital O -- as in "Origami" -- followed by zero) to the gcc compiler. In order to compile with full optimization, pass `-O3` to the compiler. You can do it by modifying the `CFLAGS` in your Makefile.

1.8 Deliverables

The assembly language and C language implementations of your "llmultiply" function. Demo the application for the lab assistant. Present your conclusions with respect to the three run times.

Fill and send in the feedback questionnaire.

1.9 Resources

NASM reference: <http://www.nasm.us/doc/>

NASM tutorial: <http://www.grack.com/downloads/djgpp/nasm/djgppnasm.txt>

Instruction set: <http://courses.ece.uiuc.edu/ece390/books/labmanual/inst-ref-general.html>

Others:

<http://www.intel.com/content/www/us/en/processors/architectures-software-developer-manuals.html>

https://en.wikipedia.org/wiki/X86_instruction_listings