

ME 333: Introduction to Mechatronics

Assignment 3: Investigating the PIC C32 Tool Suite

Electronic submission due **before** 11:00 AM on Thursday February 2nd

All questions asked in this problem set must be typed and submitted via Blackboard. See the ME333 homepage for more details. To make it easier to find the questions in this document, **questions are in bold**.

1 Understanding the Compiler and the Includes

The questions in this section are about “Simple_pic.c”, which we discussed in class. This program can be found on the Mechatronics wiki at http://hades.mech.northwestern.edu/index.php/Image:Simple_pic.c. Before answering the problems in this section, you should create an MPLAB X project with this file and compile it.

As mentioned in the *Crash Course in C* document, there are several steps in the C compilation process that takes a piece of source code in C all of the way to a .hex file that you can put on your PIC. As with any C compilation process, the main components in the mplabc32 compiler suite utilized by MPLAB X are the compiler the assembler, and the linker. In this section, we are going to look at each of components in more detail.

1.1 The C compilation process

The job of the C compiler¹ is to translate your C source code to the assembly language of the target device. The compiler has the freedom to do this in several ways and compiler options can produce significantly different assembly-level code. The assembler performs the task of taking the assembly language files (*.s file) created by the compiler and outputs object files (*.o file). The object files are the machine code equivalent of your C source files². In the final step, the linker takes all of the object files and outputs the executable file (*.elf file). The role of the linker is to put all of the object files into one file and assign memory addresses to each function and global variable defined in your C program according to a set of rules in a linker script (*.ld file, such as procdefs.ld). The process up to this point is very similar to what happens when we compile C code for the computer. When compiling code for the PIC there is an additional step that involves converting the *.elf executable into a *.hex file that we then load onto the PIC.

Questions

- (a) **What files were generated after you compiled the “Simple_pic.c” project?** This includes any *.o, *.o.d, *.elf, and *.hex files. This does not include any Makefiles, *.mk, *.bash, *.xml or *.properties files.

¹Technically the Microchip compiler is a cross-compiler, because the “host” machine (e.g., your PC or laptop) has a different hardware architecture (Intel, AMD, ARM, and MIPS are examples of companies that sell computer chips with different architectures) than the “target” machine (the PIC32).

²In fact, an object file can stand in as a replacement for the corresponding source code. If I wanted to compile my program with your proprietary functions, you can send me your object files and I would be able to compile my program without your source code.

(q1) Simple-pic.o, Simple-pic.o.d, Simple.X.production.elf, Simple.X.production.hex

- (b) In MPLAB X, click *Help*→*Help Contents* to bring up the documentation for both MPLAB X and the compiler. In the “Contents” tab you will find a topic titled “C32 Toolchain” under the “Language Tools” heading. After reading the first couple of pages³, **what are the names of tools used for compiling, assembling and linking?** These are the executables that the IDE is ultimately calling when building your program.

(q1) On p. 2 of the MPLAB X IDE and Tools Installation subsection, we find the names of the tools. The compiler is pic32-gcc, the assembler is pic32-as, and the linker is pic32-ld. If its a Windows machine, then these files end in .exe. If we wanted to, we could call these programs directly to build our C code.

- (c) In class you explored <plib.h>, which turned out to be a convenient way to include all of the Microchip PIC32 definitions. Go to the top-level Microchip directory on your computer (on my PC it's C:\Program Files\Microchip) and search for files named <plib.h> in that directory. **How many files named <plib.h> do you find?** If you open them, you'll notice that they each have subtle differences. How does the C compiler know which file to use? How can you even tell what the compiler is doing? A handy compiler switch is `--verbose`, which tells the gcc compiler to tell you everything it is doing. To enable this switch, go to *Run*→*Set Project Configuration*→*Customize....* In the left window pane, select pic32-gcc. This changes the right window pane where you should see a textbox labeled “Additional options:”. In this box type in `--verbose`. Press OK and recompile your project. After compiling your project look through the output window and you'll find the `#include` search path. Using this information, **which plib.h file is the compiler using?** Write down enough of the file's path name to distinguish it from the other file locations. If in doubt, give the absolute path.

(q1) There are 4 versions of plib.h in:

- mplabc32\v2.02\pic32-libs\include
- mplabc32\v2.02\pic32-libs\include\lega-c
- mplabc32\v2.02\pic32mx\include,
- and mplabc32\v2.02\pic32mx\include\lega-c.

(q2) My system includes the one located in the following directory

Unix/Linux: (MICROCHIP-TOP-LEVEL)/mplabc32/v2.02/pic32mx/include/plib.h

Windows: (MICROCHIP-TOP-LEVEL)\mplabc32\v2.02\pic32mx\include\plib.h

- (d) Eventually everything should map down to the hardware. Open up the p32mx795f5121.h file and go to line 5,924. You'll see the definition for PORTA. Judging from the attribute tag, **where is the header file telling the compiler to place PORTA?** When all of the source code has been compiled into object files it becomes the job of the linker to make sure that PORTA resides at the correct address in memory. All of the SFRs, including PORTA, are defined in a special object file, `processor.o`, that contains all of the SFR addresses. Our PIC's processor.o file is in pic32mx\lib\proc\32MX795F512L. **What else is in this folder that might be useful to the linker?** If you have some experience with the command line you can inspect this object file with a utility called objdump:

```
./pic32-objdump -t processor.o
```

Some sample output from executing this command is shown below:

```
bf886000 g      0 *ABS*      00000000 TRISA
bf886004 g      0 *ABS*      00000000 TRISACLR
```

³Documentation is also available in *.pdf form. It is available on Microchip's website or from within the directories where MPLAB X and your C32 compiler are installed.

bf886008	g	0	*ABS*	00000000	TRISASET
bf88600c	g	0	*ABS*	00000000	TRISAINV
bf886010	g	0	*ABS*	00000000	PORTA
bf886014	g	0	*ABS*	00000000	PORTACLR
bf886018	g	0	*ABS*	00000000	PORTASET
bf88601c	g	0	*ABS*	00000000	PORTAINV
bf886020	g	0	*ABS*	00000000	LATA
bf886024	g	0	*ABS*	00000000	LATACLR
bf886028	g	0	*ABS*	00000000	LATASET
bf88602c	g	0	*ABS*	00000000	LATAINV

At what address (the first column) is processor.o telling the linker to place PORTA? Is the address a virtual or physical address? This address should line up with the address given in the PIC32 data sheet. You can verify this by looking at Table 4-24 in the data sheet. A useful piece of information to know is the default value of PORTA at startup. Refer to Table 4-24, **what is the reset value of PORTA?**

- (q1) The header is instructing the compiler to place PORTA in the SFR memory region.
 - (q2) There is a procdefs.ld file that the linker can use. In fact, this is the default procdefs.ld file the linker uses when we don't have a procdefs in our local project directory.
 - (q3) The address is 0xBF886010.
 - (q4) The address is a virtual address.
 - (q5) The reset address is a series of x's, which means that the value is unknown on reset.
- (e) We will sparingly use the functions and macros defined in the Microchip library, but you should still be aware of what it offers. In your PIC32 C tool suite directory, you'll find a header file for each important peripheral on the PIC32 in the peripheral/ folder. Open the file ports.h in the peripheral directory and go to line 1,171. Here you'll find macros for manipulating PORTA. **How would you use mPORTAReadBits(...) to read the values on pins A0 and A9?** Your answer should involve calling the macro function with the proper input.

- (q1) Using the macros defined in ports.h:

```
unsigned int vals = 0;
vals = mPORTAReadBits(BIT_0 | BIT_9);
```

Without the macros:

```
unsigned int vals = 0;
vals = mPORTAReadBits(0x0201);
```

1.2 The Linker

Recall that it is the job of the linker to stitch the object files into a .elf file. Passing the option “-M” to the linker tells the linker to generate a “map” file. A *.map file is a simple text file that contains a report on the output of the linking process. It details how much memory was used as well as the address assigned to your program's global variables, functions, and object files. To tell MPLAB X to generate this file, click *File→Project Properties* and then select the pic32-ld section of the options. From the *Option categories* dropdown menu, select *Diagnostics*; then in the box to the right of “Generate map file” enter a name for the map file to be saved as. For example, “simple_pic.memory.map”.

Note that the *.map file is organized following a standard convention. Briefly, the .text section is where the machine instructions get grouped together. The other sections like .sbss and .data, and .bss are where constants and global data is stored.

Questions

- (a) By searching the generated *.map file, determine the following: **what memory addresses (in hex) did the linker place the main() and delay() functions in the output .elf file?** (Use a text editor with searching capabilities, you shouldn't read the whole *.map file.)

(q1) Based on our map file, the memory location of main and delay are 0x9d008008 and 0x9d008084, respectively. Your map file might differ slightly, but as long as the address is in the memory region defined by the kseg0_program_mem in the ME 333 version of the procdefs.ld file, that is between 0x9D008000 and 0x9D080000, your answer is probably fine.

- (b) Add a global variable and a global constant to your file by adding the following line just below the `#include <plib.h>` line:

```
int glob = 21;
```

Now clean and rebuild the project and open the *.map file again. Find the reference to "glob" in the *.map file. It should be in a section called ".sdata". **How many bytes does this variable use (in hex)?** Change the data type to a `char`. **How many bytes does it use now (in hex)?** Look at the "Microchip PIC32 Memory-Usage Report" towards the top of the *.map file, and find the .sdata section. **Where in the PIC's memory does char glob get placed (in hex)?**

(q1) When glob is an int, it takes 4 bytes

(q2) When glob is a char, it takes 1 byte

(q3) glob goes in kseg1 data memory at address 0xa0000000

- (c) **How large is the total .o (in hex bytes) produced by this simple C file?** This will be found in a .text section. The 32 bit hex number is the address in the virtual memory map where the object file will be placed. The shorter hex number in the third column is the total memory usage of that object file.

(q1) Based on our map file, the memory location of main and delay are 0x9d008008 and 0x9d008084, respectively. Your map file might differ slightly, but as long as the address is in the memory region defined by the kseg0_program_mem in the ME 333 version of the procdefs.ld file, that is between 0x9D008000 and 0x9D080000, your answer is probably fine.

(q2) Our map file reports that the size of the Simple_pic.o file should be between 0x104 and 0x150. The object file size depends on how many instructions are in the file. If we compile a fresh copy of Simple_pic.c downloaded from the wiki, we get 0x104, but if we uncomment every line of commented C code, we get 0x150. Depending on how you compiled your code, you might get a different value, but your file size should be near that range.

- (d) The procdefs.ld file is used by the linker to determine where in memory the object files go. **Is the location of the "Simple_pic.c" consistent with procdefs.ld, in other words, does the location of the .o get placed within kseg0_program_memory as we would expect? Are the instructions here cacheable?**

(q1) Yes, procdefs.ld specifies the start of the kseg0_program_memory as 0x9D008000. The map file puts Simple_pic.c at 0x9d008008.

(q2) The instructions here are cacheable.

- (e) Temporarily move `procdefs.ld` out of the directory that contains the source code. Clean and build the project to re-compile everything. Look at the new map file. **Where does “Simple_pic.o” get placed now (in hex)? Has it changed?** Note that the NU32 bootloader also has code in this region and it won’t erase itself. Knowing this, **will the newly compiled version (that is missing `procdefs.ld`)⁴ run on the NU32?** Having control of where programs are placed in memory according to a linker file is a powerful tool that allows different programs, like a bootloader and your code to be placed in flash without interfering with each other.

(q1) The assembly instructions contained in the `Simple_pic` object file (i.e., the code comprising the main and delay functions) are now located at `0x9d000008`.

(q2) The memory address has changed from `0x9d008008` to `0x9d000008`.

(q3) No, part of the bootloader occupies this address and the bootloader won’t overwrite itself. The end result being that this compiled version of `Simple_pic.c` is not copied onto the PIC and will not run.

1.3 The Disassembler

It is very easy to go from an object file to the original assembly file because a computer’s assembly language is a direct one-to-one mapping to the computer’s machine code. Going from a `.o` file to a `.s` file requires a disassembler –the MPLAB toolsuite uses `pic32-objdump`. If we have access to the source code, the disassembly will show C code followed by the assembly language it generated. The format of each line that contains assembly is a memory address in hex, machine code in hex, followed by the code in assembly language. To see the disassembly in MPLAB X, click *Window→Output→Disassembly listing window*. Note that a file containing the disassembly can also be generated by passing command-line arguments to the compiler, or calling `pic32-objdump` manually.

Questions

- (a) In the `delay()` function there is the following line:

```
LATABits.LATA5 = !LATABits.LATA5; LATABits.LATA4 = !LATABits.LATA4;
```

How many assembly instructions does this line generate when we view the disassembly? Just below this line is a different version of the same thing that uses the “fast bit manipulation” INV register for Latch A. Change your code to use this line instead, recompile, and analyze the new disassembly. **How many instructions are produced now?**

(q1) In its original form, it produces 18 assembly instructions

(q2) In the new form, it produces 3 assembly instructions, thus this form executes much more quickly.

- (b) Assume that you have the value `0xF0CA` in LATA. You are trying to get the latch to contain the value `0xAAAA`. **What are the two lines of C code that write to the `LATASET` and `LATACLK` addresses to accomplish this? What value do you write to the `LATAINV` address to accomplish the same thing?**

⁴In the absence of a local `procdefs.ld` file, the linker automatically uses a default one that is included with the installation of the C32 toolsuite.

(q1) LATASET = 0xA20; LATACLR = 0x5040;

(q2) LATAINV = 0x5A60;

- (c) MPLAB X has some very useful tools for exploring the files that your project includes. Clicking *Window*→*Classes* from the top menu will bring up a listing of all classes, functions, structs, macros, and constants that are available for your use. As described in detail at http://hades.mech.northwestern.edu/index.php/NU32:_A_Detailed_Look_at_Programming_the_PIC32_on_the_NU32 the inclusion of `plib.h` at the top of this simple file actually includes many `.c`, `.h`, and `.o` files written by Microchip that define all of the SFRs, the addresses of the SFRs, the macros to interface with the PIC, and much more. The *Classes* window provides a convenient interface to all of these files. Using this window, **what file defines the struct we are using when we type `LATAbits.LATA5`**?

(q1) The file is called `p32mx795f512l.h`. It is located in `mplabc32/pic32mx/include/proc/` directory.

- (d) We can roughly estimate the time that a given function or chunk of code will take to run by analyzing the assembly code. The architecture of the PIC32 is such that most instructions can be executed in a single clock cycle. Modify the for loop in the delay function, so that it stops after 10,000 iterations. Looking at the disassembly for the `delay()` function, find the line that looks something like:

```
9D0080B0  AFC00000  SW ZERO, 0(S8)
```

located just below the C line `for (j=0;j<10000;j++)`. Now find the line

```
9D0080D0  00000000  NOP
```

just above the “}” that closes the `for` loop. The memory addresses in the leftmost column may differ in your file, but the instructions should be the same. **How many lines of assembly code are there between these two lines?** Each of these assembly lines gets executed once for a single pass through the `for` loop. Knowing that our CPU runs at 80MHz, **how long does a single pass through the for loop take?** Using this number, **at approximately what frequency should the lights be toggling?** Looking at the running program, **does your estimate seem reasonable?**

- (q1) We counted 14 lines based on our disassembly file (if you included the SW and NOP instruction and counted 16 that's fine), which gave us the following lines of assembly code for the `for` loop in `delay()`:

```
25:                                for (j=0;j<10000;j++) {
9D0000FC  AFC00000  SW ZERO, 0(S8)
9D000100  0B40004B  J 0x9D00012C
9D000104  00000000  NOP
9D000120  8FC20000  LW V0, 0(S8)
9D000124  24420001  ADDIU V0, V0, 1
9D000128  AFC20000  SW V0, 0(S8)
9D00012C  8FC20000  LW V0, 0(S8)
9D000130  28422710  SLTI V0, V0, 10000
9D000134  1440FFF4  BNE V0, ZERO, 0x9D000108
9D000138  00000000  NOP
26:                                while(!PORTCbits.RC13); // Pin C13 is the USER switch
9D000108  00000000  NOP
9D00010C  3C02BF88  LUI V0, -16504
9D000110  8C426090  LW V0, 24720(V0)
```

```

9D000114 30422000 ANDI V0, V0, 8192
9D000118 1040FFFC BEQ V0, ZERO, 0x9D00010C
9D00011C 00000000 NOP
27:                                     }

```

- (q2) $14 \text{ instructions} \times 12.5 \text{ ns per instruction} = 175 \text{ ns per iteration.}$
 - (q3) The loop runs 10,000 times before exiting, so the delay function is approximately 10,000 iterations $\times 175 \text{ ns per iteration} = 1.75 \text{ ms.}$ This represents the amount of time the light is on or off, hence the light is toggled at about 571 Hz.
 - (q4) At 571 Hz, our eyes will not pick up the on/off transitions, so we expect the lights to appear on to us. When running the program the lights appear to be on, but we know that the lights should be blinking, so the rate is reasonable.
- (e) Modify the C code in this simple program to make pin A0 a digital output. Then add lines above and below the call to `delay()` that allow you to toggle the logic level that pin A0 has. If you haven't yet, modify the for loop in the delay function, so that it stops after 10,000 iterations. Now with your NUScope, calculate the time that the delay loop takes. This is a useful trick to time how long your functions are taking. **How long is delay? How far off was your estimate?**

- (q1) The `delay()` function runs in approximately 2 ms.
- (q2) The estimate was pretty good being off by 1/4 ms.

1.4 The Bootloader

When the PIC32 is reset it automatically starts executing whatever it finds at memory address 0xBFC00000. For us this happens to be the start of the NU32 bootloader program. Download the bootloader source code from http://hades.mech.northwestern.edu/index.php/Image:NU32_boot.zip. In MPLAB X load the project in the PIC32_UART_Bootloader_BootFlash_Explorer16.X project folder. The code in this project is complicated, but `BootLoader.c` is simple enough to follow along at a high level. For example, you can tell that lines 49-54 are where the configuration bits are set. The main program calls two functions, `ValidAppPresent()` and `JumpToApp()`. Take a look at these functions and answer the questions below.

Questions

- (a) In the function `JumpToApp()`, **what address is the application expected to begin at?** You should use the tools built into the IDE to quickly find the definition of `USER_APP_RESET_ADDRESS`. For example, *Navigate* \rightarrow *Go To Declaration*.
 - (q1) The `BootLoader.h` header file defines `USER_APP_RESET_ADDRESS` as 0xBD006100, so the bootloader expects the program to begin at 0xBD006100.
- (b) **How does the bootloader know in `ValidAppPresent()` that there is a valid application at `USER_APP_RESET_ADDRESS`?** Because the PIC32 uses NAND flash memory, a long string of this value represents unused space on a NAND-based flash device, like the memory card found in cameras or USB sticks. Note: If you are unsure of what a DWORD is, use the IDE to help you find its definition.
 - (q1) If no program is present, then the memory address will have the value 0xFFFFFFFF. When flash is erased the erased memory region is filled with the hex value 0xF, so the PIC assumes that if it reads a 32-bit value of non-0xF hex values, then there must be a program present.

2 What to turn in

For the questions, you must submit typed responses. Place your typed responses (the TAs prefer a PDF file) in a **zip** file and submit the zip file through Blackboard before class on the date the assignment is due. The name of the zip file you submit will be of the form lastname_firstname_a3.zip.