

15-213/18-243, Fall 2010
Lab Assignment L1: Manipulating Bits
Assigned: Aug. 24, Due: Tuesday, Sept. 7, 11:59PM
Last Possible Time to Turn in: Friday, Sept. 10, 11:59PM

Important: There is a lot of information in this handout. Be sure to study it carefully, and to reread it when you have questions before asking for help from one of the course staff.

Please contact the 15-213 staff list (15-213-staff@cs.cmu.edu) for questions about assignment.

1 Introduction

The purpose of this assignment is to become more familiar with bit-level representations of common patterns, integers, and floating-point numbers. You'll do this by solving a series of programming "puzzles." Many of these puzzles are quite artificial, but you'll find yourself thinking much more about bits in working your way through them.

1.1 Logistics

This is an individual project. All handins are electronic. Clarifications and corrections will be posted on the Autolab message board.

1.2 Logging in to Autolab

All 15-213 labs are being offered this term through a Web service developed by Hunter Pitelka and Prof. David O'Hallaron called *Autolab*. Before you can download your lab materials, you will need to update your Autolab account. Point your browser at the Autolab front page

<http://autolab.cs.cmu.edu>

You will be asked to authenticate via WebISO. After you authenticate this first time, Autolab will prompt you to update your account information with a "nickname." Your nickname is the external name that identifies you on the public scoreboards that Autolab maintains for each assignment, so pick something interesting!

You can change your nickname as often as you like. Once you have updated your account information, click on “Edit”, and then select the “Home” link to proceed to the main Autolab page.

If you added the class late, you might not be included in Autolab’s list of valid students. In this case you won’t see the 15-213/18-243 course listed on your Autolab home page. If this happens, just send email to the staff mailing list (15-213-staff@cs.cmu.edu) requesting an Autolab account, and we will add you to the list.

1.3 Obtaining your Lab Materials

Your lab materials are contained in a Unix tar file called `datalab-handout.tar`, which you can download from Autolab. After logging in to Autolab at

```
http://autolab.cs.cmu.edu
```

you can retrieve the `datalab-handout.tar` file by selecting “Datalab -> Download Lab Materials” and then hitting the “Save File” button.

Start by copying `datalab-handout.tar` to a directory in which you plan to do your work. Then give the command “`tar xvf datalab-handout.tar`”. This will create a directory called `datalab-handout` that contains a number of files. The only file you will be modifying and handing in is `bits.c`.

WARNING: Do not let the Windows WinZip program open up your `.tar` file (many Web browsers are set to do this automatically). Instead, save the file to your AFS directory and use the Linux `tar` program to extract the files. In general, for this class you should NEVER use any platform other than Linux to modify your files, doing so can cause loss of data (and important work!).

The file `btest.c` contains code that performs a simple, non-exhaustive check of the functional correctness of your code. The file `README` contains additional documentation about `BTEST`. Use the command `make` to generate the test code and run it with the command `./btest`.

The included program `DLC` can be used to check your solutions for compliance with the coding rules. The included programs `ISHOW` and `FSHOW` can be used to help examine the bit representations of integer and floating point numbers.

The files in the subdirectory `bddcheck` implement the BDD checker, a tool that formally verifies your code.

The file `driver.pl` is the main *driver program* that uses `DLC` and the BDD checker to check your `bits.c` for correctness, grades your submission, and optionally streams the results to the Autolab server for inclusion on the Autolab scoreboard. You can run the driver as often as like you to check your work.

The remaining files are used to build the program `BTEST`.

1.4 Your Assignment

The `bits.c` file contains a skeleton for each of the 13 programming puzzles. Your assignment is to complete each function skeleton according to a strict set of programming rules, intended to help you understand how values are represented at the bit-level and how to manipulate bit patterns using standard C operations.

2 Evaluation

Your code will be compiled with GCC and exhaustively tested with the BDD checker. Your score will be computed out of a maximum of 66 points based on the following distribution:

35 Correctness of code.

26 Performance of code, based on number of operators used in each function.

5 Style points, based on your instructor's subjective evaluation of the quality of your solutions and your comments.

The 13 puzzles you must solve have been given a difficulty rating between 1 and 4, such that their weighted sum totals to 35.

The code in BTEST simply tests your functions for a number of different cases. For most functions, the number of possible argument combinations far exceeds what could be tested exhaustively. To provide complete coverage, we have created an experimental *formal verification* program, CBIT that, in effect, tests your functions for all possible combinations of arguments. It does this by viewing each bit of the function result as a Boolean function of the bits comprising the function arguments. It uses a data structure known as *Binary Decision Diagrams* (BDDs) (R. E. Bryant, *IEEE Transactions on Computers*, August, 1986) to represent these Boolean functions in a way such that the program can efficiently compare the results of your functions with those of a set of reference solutions. If the bit-level functions match, then the two C functions compute identical results. Otherwise, CBIT can generate a *counterexample*, i.e., a set of function arguments where your function will produce a different result than the reference solution.

You do not invoke CBIT directly. Instead, there is a series of Perl scripts that set up and evaluate the calls to it. Execute

```
unix> ./bddcheck/check.pl -f fun
```

to check function fun. Execute

```
unix> ./bddcheck/check.pl
```

to check all of your functions.

Note: The Perl scripts are a bit picky about the formatting of your code. They expect the function to open with a line of the form:

```
int fun (...)
```

or

```
unsigned fun (...)
```

and to end with a single right brace in the leftmost column. That should be the only right brace in the leftmost column of your function.

You will get full credit for a puzzle if the BDD checker determines that your solution is correct, and no credit otherwise. The formal verification provided by the BDD checker will show you that there are no bugs lurking in your code. You'll find yourself wishing you could do this kind of testing with every program you write. Unfortunately, BDDs can handle only relatively simple functions such as the ones you are writing for this assignment. Beyond this, the BDDs get too large to represent and manipulate.

Regarding performance, our main concern at this point in the course is that you can get the right answer. However, we want to instill in you a sense of keeping things as short and simple as you can. Furthermore, some of the puzzles can be solved by brute force, but we want you to be more clever. Thus, for each function we've established a maximum number of operators that you are allowed to use for each function. Assignment operators ('=') aren't counted. This limit is very generous and is designed only to catch egregiously inefficient solutions. You will receive two points for each function that satisfies the operator limit.

Finally, we've reserved 5 points for a subjective evaluation of the style of your solutions and your commenting. Your solutions should be as clean and straightforward as possible. Your comments should be informative, describing the strategy behind your solution.

3 Bit and Integer Manipulations

The first series of puzzles involve creating common bit patterns and manipulating two's complement representations of integers. These puzzles have the strictest set of programming rules. You may only use *straightline* code (i.e., no loops or conditionals) and a limited number of C arithmetic and logical operators. Specifically, you are *only* allowed to use the following operators:

= ! ~ & ^ | + << >>

A few of the functions further restrict this list. Also, you are only allowed to use constant values between 0 and 255 (0x0 to 0xFF). See the comments in `bits.c` for detailed rules and a discussion of the desired coding style.

For the bit and integer manipulation puzzles, you may not use any control structures such as loops, function calls, and conditionals within your code. You also may not do any casting or use any data types other than `int`. You may not use any unions, structs, or arrays. You may assume that data type `int` is 32 bits long and encodes integers in two's complement format. Both left and right shifts require a shift amount between 0 and 31, and right shifts are performed arithmetically. **WARNING:** Do not try to shift by more than 31, bad things will happen! (See Section 4 for coding rules for other puzzles.)

3.1 Part I: Bit Manipulations

Table 1 describes a set of functions that manipulate and test sets of bits. The "Rating" field gives the difficulty rating (the number of points) for the puzzle, and the "Max ops" field gives the maximum number of operators you are allowed to use to implement each function. See the comments in `bits.c` for more

Name	Description	Rating	Max Ops
<code>implication(x,y)</code>	$x \rightarrow y$	2	5
<code>byteSwap(x,n,m)</code>	Swaps nth byte in x with mth byte in x.	2	25
<code>conditional(x,y,z)</code>	Same as $x ? y : z$	3	16
<code>bang(x)</code>	Compute $!x$ without using $!$	4	12
<code>isAsciiDigit(x)</code>	Does x represent an ASCII digit?	3	15
<code>allEvenBits(x)</code>	Are all even-numbered bits in x set to 1?	2	12
<code>logicalShift(x,n)</code>	Shift x to the right n bits.	3	20

Table 1: Bit-Level Manipulation Functions.

Name	Description	Rating	Max Ops
<code>tmin(x)</code>	Return the minimum two's complement integer	1	4
<code>leastBitPos(x)</code>	Return mask of least significant one bit	2	6
<code>isLessOrEqual(x,y)</code>	Return $x \leq y$	3	24
<code>howManyBits(x)</code>	Returns the minimum number of bits needed to represent x.	4	90

Table 2: Arithmetic Functions

details on the desired behavior of the functions. You may also refer to the test functions in `tests.c`. These are used as reference functions to express the correct behavior of your functions, although they don't satisfy the coding rules for your functions.

3.2 Part II: Two's Complement Arithmetic

Table 2 describes a set of functions that make use of the two's complement representation of integers. Again, refer to the comments in `bits.c` and the reference versions in `tests.c` for more information.

4 Floating-Point Operations

For this part of the assignment, you will implement some common single-precision floating-point operations. In this section, you are allowed to use standard control structures (conditionals, loops), and you may use both `int` and `unsigned` data types, including arbitrary unsigned and integer constants. You may not use any unions, structs, or arrays. Most significantly, you may not use any floating point data types, operations, or constants. Instead, any floating-point operand will be passed to the function as having type `unsigned`, and any returned floating-point value will be of type `unsigned`. Your code should perform the bit manipulations that implement the specified floating point operations.

Name	Description	Rating	Max Ops
<code>float_neg(f)</code>	<code>-f</code>	2	10
<code>float_f2i(f)</code>	<code>(int) f</code>	4	30

Table 3: Floating-Point Functions. Value `f` is the floating-point number having the same bit representation as the unsigned integer `uf`.

4.1 Part III: Floating-Point Arithmetic

Table 3 describes a set of functions that operate on the bit-level representations of floating-point numbers. Refer to the comments in `bits.c` and the reference versions in `tests.c` for more information.

The included program `FSHOW` helps you understand the structure of floating point numbers, letting you decipher the results from the BDD checker on floating-point problems. When the BDD checker finds a counterexample, it prints the decimal value of argument(s) that cause a discrepancy between the puzzle code and the reference version. To see what these bit patterns represent as a floating-point number, use `FSHOW`, e.g.:

```
unix> ./fshow 2080374784

Floating point value 2.658455992e+36
Bit Representation 0x7c000000, sign = 0, exponent = f8, fraction = 000000
Normalized. 1.0000000000 X 2^(121)
```

You can also give `FSHOW` hexadecimal and floating point values, and it will decipher their bit structure.

Functions `float_abs` and `float_half` must handle the full range of possible argument values, including not-a-number (NaN) and infinity. The IEEE standard does not specify precisely how to handle NaN's, and the IA32 behavior is a bit obscure. We will follow a convention that when of these functions is given a NaN value as argument, it returns the same value for the result.

5 Advice

You can work on this assignment using one of the class shark machines or one of the the Andrew Linux servers (`ssh unix.andrew.cmu.edu`). The BDD checker and the DLC program are distributed as 32-bit Linux executables, and so you'll need to be working on a compatible Linux machine in order to use those tools. In general, we recommend that you work on the shark machines, since these are the systems you'll be using for most of the labs in this course.

The DLC program, a modified version of an ANSI C compiler, will be used to check your programs for compliance with the coding style rules. The typical usage is

```
unix> ./dlc bits.c
```

Type `./dlc -help` for a list of command line options. The `README` file is also helpful.

- The DLC program runs silently unless it detects a problem.
- Andrew Linux machines have a program called `/usr/local/bin/dlc`, which is *not* the same as our DLC program. So always run DLC using a full path name:

```
unix> ./dlc bits.c
```

- Don't include the `<stdio.h>` header file in your `bits.c` file, as it confuses DLC and results in some non-intuitive error messages. You will still be able to use `printf` in your `bits.c` file for debugging without including the `<stdio.h>` header, although GCC will print a warning that you can ignore.
- The DLC program enforces a stricter form of declarations than is the case for C++ or Java or even that is enforced by GCC. In particular, any declaration must appear in a block (what you enclose in curly braces) before any statement that is not a declaration. For example, it will complain about the following code:

```
int foo(int x)
{
    int a = x;
    a *= 3;    /* This statement is not a declaration */
    int b = a; /* ERROR: Declaration not allowed here */
}
```

The BDD checker cannot handle functions that call other functions, including `printf`. You should use BTEST to evaluate code with debugging `printf` statements. Be sure to remove any of these debugging statements before handing in your code.

Check the file README for documentation on running the BTEST program. You'll find it helpful to work through the functions one at a time, testing each one as you go. You can use the `-f` flag to instruct BTEST to test only a single function, e.g., `./btest -f bitAnd`. You can feed it specific function arguments using the option flags `-1`, `-2`, and `-3`. Also, the `-g` option is a nice way to get a compact summary of the correctness of each function.

The testing provided by BTEST is somewhat weak for the floating-point problems, where there are tricky issues of denormalized numbers, rounding, and overflow. Use the BDD checker to detect problems in your code, but then use BTEST and FSHOW (or ISHOW for integer problems) to help you better understand what's going on.

6 Submission Instructions

Unlike other courses you may have taken in the past, in this course you may handin your work as often as you like until the due date of the lab. There are two types of handins: *unofficial* and *official* handins.

- **Unofficial handins.** As you work on the assignment you can use the driver program `driver.pl` to stream your current results to the Autolab server to be displayed on the public real-time scoreboard

that Autolab maintains for this assignment. The driver is the same program our autograder calls when it grades your handin. If your userid is bovik, then typing

```
unix> ./driver.pl -u bovik
```

will stream your results (in the form of an ASCII text line we call an *autoresult string*) to the Autolab server. The autoresult strings are logged, and the last autoresults from each student are periodically summarized on the scoreboard, under each student's Autolab nickname.

The Autolab Data Lab page provide an option that allows you to view the scoreboard (“View Class Scoreboard”). The scoreboard is updated every few seconds.

- **Official handins.** The autoresult strings sent from your local copy of the driver program are unofficial and just for fun. To receive credit, you will need to upload your `bits.c` file using the Autolab option “Handin your work for credit”. Each time you handin your code, the server will run the driver program on your handin file and produce a grade report (it also logs an official autoresult string for the class status page). The server archives each of your submissions and resulting grade reports, which you can view anytime using the “View your handin history and scores” option.

Notes:

- At any point in time, your most recently uploaded file is your official handin. You may handin as often as you like.
- Each time you handin, you should use the “View your handin history and scores” option to confirm that your handin was properly autograded.
- You must remove any extraneous print statements from your `bits.c` file before handing in.
- If you have any questions about this lab, the Autolab system, or the course in general, please email the course staff at `15-213-staff@cs.cmu.edu`. We respond to emails 24/7 and are very good about getting back to you fast. Remember: We're here to help!