

Computer Systems
baiCOSY06, Fall 2016
Data Lab: Manipulating Bits
Assigned: Sept. 5, Due: Wed., Sept. 14, 11:59.

Arnoud Visser (A.Visser@uva.nl) is the lead person for this assignment. Giulio Stramondo, Janosh Haber, Kyrian Maat, Tycho van der Ouderaa, Joop Pascha and Thomas Schaper will be the teaching assistants.

1 Introduction

The purpose of this assignment is to become more familiar with bit-level representations of 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.

2 Logistics

For this assignment you are requested for pairs. All handins are electronic. As part of this assignment, each pair is requested to keep a labjournal. See section 6 for more details. Clarifications and corrections will be posted on the course Web page.

3 Linux workstations

At the Universiteit van Amsterdam most rooms are reserved for laptop usage. The few machines which are still available at the Study Center are double boot. When the machine is running Windows, you can reach the Linux environment by restarting the machine and pressing the F2-key when asked.

Yet, this course expects that you bring a laptop with you.

There are three ways to use your laptop.

3.1 student server acheron

Inside the uva-domain you could access a Linux-server dedicated for students (`acheron.fnwi.uva.nl`). To access the server `acheron`, one first has to build up an vpn-connection with your UvAID, even when you are in the fnwi-domain. The server is als always available and could be used for programming. To login from your local machine you need a program as `ssh` or `putty`.

At `acheron`, you can find the software for this course in `/opt/prac/cs_ai`.

3.2 virtual machine

In `datanose` you find under Tools→Software a list of packages (<https://datanose.nl/#byod>). One of those is VMWare, select Player 7 (Linux or Windows) or Fusion (Mac).

Within the player, you open the Virtual Machine file that you have downloaded from: <http://sbt.science.uva.nl/bterwijn/KI2016/>.

3.3 native Linux machine

Install with `sudo apt-get install` the following packages:

- `valgrind` `kcachegrind`

4 Handout Instructions

Each group of students will get a personalized puzzle. To obtain your group's puzzle, one (and only one) of the group members get one from the puzzles. When on `acheron`, this could be done with the following commands (replace ?? with your favorite number between 1 and 99):

```
mkdir ./datalab
mkdir ./datalab/puzzle??
cd ./datalab/puzzle??
cp -pr /opt/prac/cs_ai/datalab/2016/handout/puzzle?? .
```

In all other cases, you can get one of the puzzles by downloading them from `acheron` with the commands (replace not only ??, but also `uvoid` with your studentnumber. Don't forget the space and the dot with the `scp`-command):

```
mkdir ./datalab
mkdir ./datalab/puzzle??
cd ./datalab/puzzle??
scp uvoid@acheron.fnwi.uva.nl:/opt/prac/cs_ai/datalab/2016/handout/puzzle?*/* .
```

The only file you will be modifying and turning in is `bits.c`.

The `bits.c` file contains a skeleton for each of the 15 programming puzzles. Your assignment is to complete each function skeleton using only *straightline* code for the integer puzzles (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 eight operators:

`! ~ & ^ | + << >>`

A few of the functions further restrict this list. Also, you are not allowed to use any constants longer than 8 bits. See the comments in `bits.c` for detailed rules and a discussion of the desired coding style.

5 The Puzzles

This section describes the puzzles that you will be solving in `bits.c`. You will solve puzzles with an increasing difficulty, it is not required that you solve them all.

5.1 Bit Manipulations

The first part of the assignment is a set of functions that manipulate and test sets of bits. See the comments in `bits.c` for more details on the desired behavior of the functions. In the comments also the “Rating” and “Max ops” fields are specified. 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. 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.

5.2 Two’s Complement Arithmetic

The intermediate part of the assignment involve 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.

5.3 Floating-Point Operations

For the last phase 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.

Functions `float_neg` and `float_twice` 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 for any function returning a NaN value, it will return the one with bit representation `0x7FC00000`.

The included program `fshow` helps you understand the structure of floating point numbers. To compile `fshow`, switch to the handout directory and type:

```
make
```

You can use `fshow` to see what an arbitrary pattern represents as a floating-point number:

```
./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.

6 Labjournal

To solve the puzzles, you need creativity, knowledge and experimental skills. It is required that you document your progress on finding a solution on the puzzles in a labjournal. See the page 'Het labboek'¹ on the website on Academic Skills for some general pointers about what a labjournal should be contain.

7 Evaluation

Your score will be computed out of a maximum of 37 points based on the following distribution:

15 Correctness points.

12 Performance points.

10 Academic style points.

Correctness points. The 6 puzzles you must solve have been given a difficulty rating between 1 and 4, such that their weighted sum totals to 16. We will evaluate your functions using the `btest` program, which is described in the next section. You will get full credit for a puzzle if it passes all of the tests performed by `btest`, and no credit otherwise.

Performance points. 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

¹<http://www.practicumav.nl/onderzoeken/labboek.html>

we've established a maximum number of operators that you are allowed to use for each function. This limit is very generous and is designed only to catch egregiously inefficient solutions. You will receive two points for each correct function that satisfies the operator limit.

Style points. Finally, your labjournal will be evaluated by your tutors on the criteria on the "labboek" page from the website on Academic Skills.

Autograding your work

We have included some autograding tools in the handout directory — `btest`, `dlc`, and `driver.pl` — to help you check the correctness of your work.

- **btest:** This program checks the functional correctness of the functions in `bits.c`. To build and use it, type the following two commands:

```
make
./btest
```

Notice that you must rebuild `btest` each time you modify your `bits.c` file.

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:

```
./btest -f bitAnd
```

You can feed it specific function arguments using the option flags `-1`, `-2`, and `-3`:

```
./btest -f bitAnd -1 7 -2 0xf
```

Check the file `README` for documentation on running the `btest` program.

- **dlc:** This is a modified version of an ANSI C compiler from the MIT CILK group that you can use to check for compliance with the coding rules for each puzzle. The typical usage is:

```
./dlc bits.c
```

The program runs silently unless it detects a problem, such as an illegal operator, too many operators, or non-straightline code in the integer puzzles. Running with the `-e` switch:

```
./dlc -e bits.c
```

causes `dlc` to print counts of the number of operators used by each function. Type `./dlc -help` for a list of command line options.

- **driver.pl:** This is a driver program that uses `btest` and `dlc` to compute the correctness and performance points for your solution. It takes no arguments:

```
./driver.pl
```

Your instructors will use `driver.pl` to evaluate your solution.

8 Handin Instructions

To handin your solution, upload your labjournal to Blackboard and mail your file `bits.c` to your assistant.

9 Advice

- Homework assignments 2.58 - 2.70 are the same style of puzzles as the one in this assignment. Also look at page 192 for the example code of problem 2.42.
- 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 C declarations than is the case for C++ or 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;    /* Statement that is not a declaration */
    int b = a; /* ERROR: Declaration not allowed here */
}
```