

CPE 323

MODULE 02

NUMBER REPRESENTATIONS AND DATA TYPES IN MODERN COMPUTERS: A REVIEW

Aleksandar Milenković

The LaCASA Laboratory, ECE Department, The University of Alabama in Huntsville

Email: milenka@uah.edu

Web: <http://www.ece.uah.edu/~milenka>

Overview

This module reviews data representations in digital computers. It starts with positional numeral systems - decimal, binary, octal, and hexadecimal, then discusses storing and interpreting numbers in digital computers, positive and negative integers and arithmetic operations on them. Finally, it discusses representations of fractional numbers in digital computers (fixed-point and floating-point).

Objectives

Upon completion of this module learners will be able to:

- *Use different numeral systems for representing data: decimal, binary, octal, hexadecimal*
- *Convert decimal numbers into their binary, octal, and hexadecimal representations*
- *Perform arithmetic operations on different representations of integers and determine condition code flags (Carry, Overflow, Zero, and Negative)*
- *Use fixed and floating-point representations of real numbers*

Contents

1	N numeral Systems: Decimal, binary, hexadecimal, and octal	3
2	Conversion to binary from other numeral systems	4
3	Storing and interpreting information in modern computers	5
4	Integers	6
4.1	Two's Complement	6
4.2	Calculating two's complement.....	7
4.3	Arithmetic Operations: Addition.....	8

4.4	Arithmetic Operations: Subtraction	9
4.5	Arithmetic Operations: Multiplication	11
4.6	Binary Coded Decimal Numbers	12
5	Characters.....	13
6	Fraction Numbers	14
6.1	Fixed-point	14
6.2	Floating-point.....	15
7	References	19
8	Exercises	20



1 Numeral Systems: Decimal, binary, hexadecimal, and octal

We ordinarily represent numbers using the positional decimal number system that has 10 as its base (also known as the base-10 or denary system). A positional numeral system is a system for the representation of numbers by an ordered set of numeral symbols (called digits) in which the value of a numeral symbol depends on its position - for example, the "ones place", "tens place", "hundreds place," and so on. Note: other notations such as Roman are not positional. The decimal numeral system uses 10 symbols (digits): 0, 1, 2, 3, 4, 5, 6, 7, 8, and 9. It is the most widely used numeral system, perhaps because humans have ten fingers over both hands.

The decimal number 456 could also be represented as four hundreds (10^2), 5 tens (10^1), and 6 ones (10^0) and can be written as follows:

position	2	1	0
weight	10^2	10^1	10^0
digit	4	5	6

$$456_{10} = 4 \cdot 10^2 + 5 \cdot 10^1 + 6 \cdot 10^0 = 400 + 50 + 6$$

The binary positional numeral system, or the base-2 (radix-2) number system, represents numeric values using only two symbols, 0 and 1. These symbols can be directly implemented using logic gates and that is why all modern computers internally use the binary system to store information.

For example, 10111_2 represents a binary number with 5 binary digits. The left-most bit is known as the most significant bit (msb) – in our case it has a bit position of 4. The rightmost bit is known as the least-significant bit (lsb) or the bit at position 0. This binary number 10111_2 can be converted to a corresponding decimal number as follows:

position	4	3	2	1	0
weight	2^4	2^3	2^2	2^1	2^0
digit	1	0	1	1	1

$$10111_2 = 1 \cdot 2^4 + 0 \cdot 2^3 + 1 \cdot 2^2 + 1 \cdot 2^1 + 1 \cdot 2^0 = 16 + 4 + 2 + 1 = 23_{10}$$

Hexadecimal (base-16, hexa, or hex) is a positional numeral system with the base 16. It uses sixteen distinct symbols. The symbols 0–9 to represent values zero to nine, and A, B, C, D, E, F (or *a* through *f*) represent values ten to fifteen, respectively.

For example, the hex number $1A3_{16}$ corresponds to:

position	2	1	0
weight	16^2	16^1	16^0
digit	1	A	3

$$1A3_{16} = 1 \cdot 16^2 + 10 \cdot 16^1 + 3 \cdot 16^0 = 256 + 160 + 3 = 419_{10}$$

To convert a hex number into its binary format, simply replace each hex symbol with its 4-bit binary counterpart (0 – 0000_2 , 1 – 0001_2 , . . . , 9 – 1001_2 , A – 1010_2 , B – 1011_2 , C – 1100_2 , D – 1101_2 , E – 1110_2 , F – 1111_2).

$1A3_{16} = 1_1010_0011_2$ (the '_' character is here just to improve readability)

Hexadecimal numbers are used as a human friendly representation of binary coded values, so they are often used in digital electronics and computer engineering. Since each hexadecimal digit represents four binary digits (bits), it is a compact and easily translated shorthand to express values in base two.

Octal (also base-8) is a number system with the base 8 that uses 8 distinct symbols 0, 1, 2, 3, 4, 5, 6, and 7. The octal number 372_8 corresponds to:

position	2	1	0
weight	8^2	8^1	8^0
digit	3	7	2

$$372_8 = 3 \cdot 8^2 + 7 \cdot 8^1 + 2 \cdot 8^0 = 192 + 56 + 2 = 250_{10}$$

2 Conversion to binary from other numeral systems

To convert a decimal positive number to its binary equivalent, the number is divided by two, and the remainder is the least-significant bit at position 0. The result is again divided by two and the remainder is the next bit at position 1. This process repeats until the result of division becomes zero.

Example 2-1. Convert 23_{10} to binary.

To convert 23 to binary, the following steps are performed:

Operation	Remainder	(Position)
$23 \div 2 = 11$	1	(bit 0)
$11 \div 2 = 5$	1	(bit 1)
$5 \div 2 = 2$	1	(bit 2)
$2 \div 2 = 1$	0	(bit 3)
$1 \div 2 = 0$	1	(bit 4)

Reading the sequence of remainders from the bottom up gives the binary numeral 10111_2 .

This method works for conversion from any base, but there are better methods for bases which are powers of two, such as octal and hexadecimal given below.

As described above, to convert a binary number to a decimal number use the following approach:

$$10111_2 = 1 \cdot 2^4 + 0 \cdot 2^3 + 1 \cdot 2^2 + 1 \cdot 2^1 + 1 \cdot 2^0 = 16 + 4 + 2 + 1 = 23_{10}$$

Octal and hexadecimal numbers are easily converted to binary numbers and vice versa. For example, an octal number 372_8 is converted into binary by simply replacing each octal digit with its binary counterpart, $011_111_010_2$. To get a corresponding hex representation, create groups of 4 bits starting from the least significant bit. By replacing each group with a corresponding hex digit, a hex representation is created (see below).

$$372_8 = 011_111_010_2 = 0FA_{16}$$

$$1011_0101_2 = 265_8 = B5_{16}$$

3 Storing and interpreting information in modern computers

An n -bit binary number can differentiate between 2^n things. In digital computers information and space are organized in bytes. A byte has 8 bits and can represent up to $2^8 = 256$ things.

Definitions:

- 1 byte = 8 binary digits = 8 bits (e.g. 1001_0011₂)
- ½ byte = 1 nibble = 4 bits
- In 16-bit computers 1 word = 2 bytes (16 bits).
- In 32-bit computers 1 word = 4 bytes (32 bits), a half-word is 2 bytes (16-bits), and a double word is 8 bytes (64 bits).

Meaning of bits and bytes is assigned by the convention. Some examples of common encoding formats are as follows:

- 1 byte ASCII = one of 256 alphanumeric or special-purpose text characters (definitions are in the ASCII table, see <http://en.wikipedia.org/wiki/ASCII>)
- 1 byte = one short unsigned integer (0 – 255)
- 1 byte = one short signed integer (-128 – 127)
- 1 byte = two Binary Coded Decimal (BCD) digits (i.e. one per nibble)
- 1 byte = two hexadecimal (hex) digits (one per nibble)
- 1 byte = eight individual bits (e.g. 8 status flags)
- 2 bytes = one unsigned integer (0 – 65,535)
- 2 bytes = one signed integer (-32,768 – 32,767)
- 4 bytes = one unsigned long integer (0 – 4,294,967,295)

For example, the IAR Software Development Environment for MSP430 assumes the following conventions for standard C data types (see Table 1).

Table 1. C data types in IAR, sizes, ranges, and alignments in memory

Data type	Size	Range	Alignment
bool	8 bits	0 to 1	1
char	8 bits	0 to 255	1
signed char	8 bits	-128 to 127	1
unsigned char	8 bits	0 to 255	1
signed short	16 bits	-32768 to 32767	2
unsigned short	16 bits	0 to 65535	2
signed int	16 bits	-32768 to 32767	2
unsigned int	16 bits	0 to 65535	2
signed long	32 bits	-2^{31} to $2^{31}-1$	2
unsigned long	32 bits	0 to $2^{32}-1$	2
signed long long	64 bits	-2^{63} to $2^{63}-1$	2
unsigned long long	64 bits	0 to $2^{64}-1$	2

float	32 bits		2
double	64 bits		2

4 Integers

The term integer is used in computer engineering/science to refer to a data type which represents some finite subset of the mathematical integers. Integers may be unsigned or signed.

An n -bit unsigned integer is defined as follows:

$$A = A_{n-1}A_{n-2} \dots A_1A_0 = A_{n-1} \cdot 2^{n-1} + A_{n-2} \cdot 2^{n-2} + \dots + A_1 \cdot 2^1 + A_0 \cdot 2^0$$

An n -bit unsigned integer can encode 2^n numbers that represent the non-negative values from 0 ($A_{n-1}A_{n-2} \dots A_1A_0 = 00 \dots 00$) through 2^n-1 ($A_{n-1}A_{n-2} \dots A_1A_0 = 11 \dots 11$).

There are three different ways to represent negative numbers in a binary numeral system. The most common is **two's complement**, which allows a signed integer type with n bits to represent numbers from $-2^{(n-1)}$ through $2^{(n-1)}-1$. Two's complement arithmetic is convenient because there is a perfect one-to-one correspondence between representations and values, and because addition, subtraction and multiplication do not need to distinguish between signed and unsigned integer types (thus a single hardware resource can be used to perform operations on both unsigned and signed integers).

The other possibilities for representing signed numbers are *sign-and-magnitude* and *ones' complement*. The *sign-and-magnitude* representation uses $n-1$ bits to convey the magnitude with the most significant bit (MSB) used for sign (0 for +, 1 for -). The problem with this approach is that there exists two representations for value 0 ($00\dots00_2$ for a positive 0 and $10\dots00_2$ for a negative 0). With *ones' complement* representation a negative binary number is created by applying the bitwise NOT to its positive counterpart. Like *sign-and-magnitude* representation, ones' complement has two representations of 0: 00000000_2 (+0) and 11111111_2 (-0). As an example, the ones' complement form of 00101011_2 (43) becomes 11010100_2 (-43). The range of signed numbers using ones' complement in a conventional eight-bit byte is -127_{10} to $+127_{10}$.

We focus on two's complement because it is the dominant way to represent signed integers today.

4.1 Two's Complement

In the two's complement notation, a positive number is represented by its ordinary binary representation, using enough bits so that the MSB bit, also known as the sign bit, is 0. An n -bit integer in two's complement is defined as follows:

$$A = A_{n-1}A_{n-2} \dots A_1A_0 = -A_{n-1} \cdot 2^{n-1} + A_{n-2} \cdot 2^{n-2} + \dots + A_1 \cdot 2^1 + A_0 \cdot 2^0$$

When the MSB bit is zero, $A_{n-1} = 0$, a positive number in range from 0

($A_{n-1}A_{n-2} \dots A_1A_0 = 00 \dots 00$) to $(2^{n-1} - 1)$ ($A_{n-1}A_{n-2} \dots A_1A_0 = 01 \dots 11$) is represented.

When the MSB bit is one, $A_{n-1} = 1$, a negative number from -2^{n-1}

($A_{n-1}A_{n-2} \dots A_1A_0 = 10 \dots 00$) to -1 ($A_{n-1}A_{n-2} \dots A_1A_0 = 11 \dots 11$) is represented. So, the range of signed numbers that can be represented with n bits in 2's complement is from -2^{n-1} to $(2^{n-1} - 1)$. There is only one representation of 0 ($00 \dots 00_2$).

Table 2 shows ranges of signed and unsigned integers depending on the number of bits used for representation.

Table 2. Common integer data types and ranges.

Bits	Name	Range
n	n-bit integer (general case)	Signed: (-2^{n-1}) to $(2^{n-1} - 1)$
		Unsigned: 0 to $(2^n - 1)$
8	byte, octet	Signed: -128 to +127
		Unsigned: 0 to +255
16	halfword, word	Signed: -32,768 to +32,767
		Unsigned: 0 to +65,535
32	word, doubleword, longword	Signed: -2,147,483,648 to +2,147,483,647
		Unsigned: 0 to +4,294,967,295
64	doubleword, longword, long long, quad, quadword	Signed: -9,223,372,036,854,775,808 to +9,223,372,036,854,775,807
		Unsigned: 0 to +18,446,744,073,709,551,615
128	octaword	Signed: -170,141,183,460,469,231,731,687,303,715,884,105,728 to +170,141,183,460,469,231,731,687,303,715,884,105,727
		Unsigned: 0 to +340,282,366,920,938,463,463,374,607,431,768,211,455

4.2 Calculating two's complement

The two's complement operation is the negation operation. In finding the two's complement of a binary number, the bits are first inverted, or "flipped", by using the bitwise NOT operation and then the constant 1 is added to that value.

Let's assume 8-bit signed binary numbers. Consider A, $A = 23_{10} = 17_{16} = 00010111_2$. We would like to find the two's complement for A or (-A). The steps are shown below.

Example 4-1. Find -A, A=23.

```

A:      0001 0111      (23)
-----
 $\bar{A}$ :    1110 1000      (first complement of A)
+                1      (1)
=====
-A:     1110 1001      (-23)

```

So the two's complement representation of $-23_{10} = 11101001_2 = E9_{16}$.

You may shorten this process by using hexadecimal representation as follows.

Example 4-2. Find $-A$, $A=23$.

```

A                :    17
one's complement of A :    E8
                    +     1
                    -----
                    E9
  
```

A shortcut to manually convert a binary number into its two's complement is to start at the least significant bit (LSB), and copy all the zeros (working from the LSB toward the MSB) until the first 1 is reached; then copy that 1, and flip all the remaining bits. This shortcut allows you to convert a number to its two's complement without first forming its ones' complement. For example: the two's complement of "00010111" is "11101001", where the underlined digits are unchanged by the copying operation.

4.3 Arithmetic Operations: Addition

Adding two's-complement numbers requires no special processing if the operands have opposite signs: the sign of the result is determined automatically. Consider the following example of finding $A + B$, where $A=15$, $B=-5$, and $n=8$. The result is $S=0000_1010$ in binary or 10 in decimal.

Example 4-3. Add $A=15$ and $B=-5$ if $n=8$ (8-bit numbers).

```

          1111 1111 (carry bits= $C_8C_7\dots C_1$ )
          -----
A:      0000 1111   (15)
B:      + 1111 1011  (-5)
=====
S:      0000 1010   (10)
  
```

This process depends upon restricting arithmetic operations to 8 bits of precision; a carry to the 8th (nonexistent) bit position, C_8 , is ignored, producing an arithmetically correct result in two's complement. The Carry (C) flag is set to the most significant carry bit, $C=C_8$. Please note that the Carry flag does not have any particular meaning for signed numbers in 2's complement. However, if 8-bit values represent unsigned numbers (15 and 251 in our example), the Carry flag indicates that the result of addition is outside the range of numbers that can be represented (from 0 to 255). And indeed, $251+15=266$, but our 8-bit result contains $266 \% 256 = 10$. Using the carry bit as the most significant field of the 9-bit result would give a correct result of the unsigned addition.

The last two bits of the carry row (C_8 and C_7) contain vital information for 2's complement arithmetic: whether the calculation resulted in an arithmetic overflow (V flag), which means that the result is outside the range of the 2's complement system. An overflow condition exists when a carry (an extra 1) is generated into but not out of the sign bit, or out of but not into the sign bit. As mentioned above, the sign bit is the leftmost bit of the result.

If the last two carry bits are both 1's or both 0's, the result is valid; if the last two carry bits are "1 0" or "0 1", a sign overflow has occurred. Conveniently, an XOR operation on these two bits can quickly determine if an overflow condition exists.

Example 4-4. Find $A + B$, if $A=7$ and $B=3$, $n=4$.

```

          0111 (carry bits)
-----
A:      0111   (7)
B:    + 0011   (3)
=====
S:      1010   (-6) Invalid in 2's complement

```

In the example above, the far left two carry bits are "01", which means there was a two's-complement addition overflow. That is, ten is outside the permitted range for 4-bit numbers in 2's complement, which ranges from -8 to 7.

An alternative way to calculate Overflow (V) flag is to consider the sign bits of the operands and the result. Let us assume we add two 8-bit integers in two's complement, A and B, and the result is an 8-bit S. In case of addition, the overflow flag occurs if one of the following two conditions are met: (a) both operands A and B are positive numbers ($A_7=0$ and $B_7=0$), and the result S is a negative number ($S_7=1$), or (b) both operands A and B are negative numbers ($A_7=1$ and $B_7=1$), and the result S is a positive number ($S_7=0$). Consequently, the overflow flag can be calculated as follows:

$$V_{Add} = \overline{A_7} \cdot \overline{B_7} \cdot S_7 + A_7 \cdot B_7 \cdot \overline{S_7}$$

In addition to the Carry flag and the Overflow flag, two more flags are commonly used in digital computers, Zero (Z) and Negative (N). The Z flag is set when the result of the current arithmetic or logic operation is equal to 0 ($S=0$) and reset otherwise. The negative flag keeps information about the sign bit of the result S ($N=S_7$).

4.4 Arithmetic Operations: Subtraction

Computers usually use the method of complements to implement subtraction for numbers in two's complement. Although using complements for subtraction is related to using complements for representing signed numbers, they are independent. Direct subtraction works with two's-complement numbers as well. Like addition, the advantage of using two's complement is the elimination of examining the signs of the operands to determine if addition or subtraction is needed. Consider the following example.

Example 4-5. Find $A - B$ if $A=15$ and $B=-5$, $n=8$.

```

R:      11110 000 (borrow bits)
-----
A:      0000 1111      (15)
B:      - 1111 1011      (-5)
=====
S:      0001 0100      (20)

```

Overflow is detected the same way as for addition, by examining the two leftmost (most significant) bits of the borrow row (R); overflow occurs if they are different. Alternatively, the overflow flag can be calculated as follows:

$$V_{Sub} = \overline{A_7} \cdot B_7 \cdot S_7 + A_7 \cdot \overline{B_7} \cdot \overline{S_7}$$

An easy way to understand this equation is to ask yourself the following: when does an overflow occur in the operation $S = A - B$. If both numbers A and B are positive ($A_7=0$ and $B_7=0$) or negative ($A_7=1$ and $B_7=1$), the overflow cannot occur because the result will be within the range we can represent. The overflow occurs if A is a positive and B is a negative and the result S turns out to be negative – since A is a positive, subtracting a negative B would result in an even greater positive result – however, if the result exceeds the range of positive numbers that can be represented, the sign bit will be set, representing a negative number in two's complement. This condition is captured by the first term in the equation above. Another scenario when overflow occurs is if A is a negative number, B is a positive number, and the result is a positive number. Again, expected result is a negative number, but if it exceeds the range of negative numbers that can be represented, the result will turn out to be a positive number. The second term in the equation above captures that scenario. If either term is computed to a logic one, the overflow bit is set.

Example 4-6. Find $A - B$ if $A=15$ and $B=35$, $n=8$.

```

R:      11100 000 (borrow bits)
-----
A:      0000 1111      (15)
B:      - 0010 0011      (35)
=====
S:      1110 1100      (-20)

```

Arithmetic, logic, and shift operations on operands are performed in Arithmetic Logic Units (ALUs). Typically an adder is used for both addition and subtraction operations. To perform a subtract operation, $S = A - B$, an equivalent operation in two's complement can be carried out as follows: $S = A + (\overline{B} + 1)$. Remember, we already showed that $-B = \overline{B} + 1$. The first-complement of B is computed by inverting the input operand B , and adding a constant one is

achieved by setting the input Carry bit to 1 ($C_0=1$). This way, subtract operations are implemented using adder logic. Often that is the case and the flags are set following the rules for additions.

Consider the previous example of finding $A - B$, if $A=15$ and $B=35$.

Example 4-7. Find $A - B$ if $A=15$ and $B=35$, $n=8$.

```

R      00011 111 (carry bits)
-----
A:     0000 1111
 $\bar{B}$ :  + 1101 1100
      +      1
=====
S:     1110 1100      (-20)

```

The result is identical, $S=1110_1100_2$ (-20). The flag bits are set as follows:
 $C=C_8=0$; $Z=0$; $V=C_8 \text{ xor } C_7=0 \text{ xor } 0=0$; $N=S_7=1$.

A careful reader would notice that the most significant bit of the borrow row $R_8=1$ is different from the most significant bit of the carry row ($C_8=0$) in the examples above. Processor manufacturers may decide on the meaning of the Carry bit for subtraction. If the carry flag, C , should keep the value that is equivalent to R_8 and the adders are used for subtractions, then simply the bit C_8 is inverted and assigned to the carry bit ($C=\overline{C_8}$) when a subtraction is performed. In other architectures the carry bit simply takes the value of C_8 as is. In this course we will always assume the latter approach.

4.5 Arithmetic Operations: Multiplication

To encode the product of two n -bit numbers we need $2n$ bits. If we sign-extend n -bit operands (multiplier and multiplicand) to $2*n$ bits before the multiplication, direct multiplication (discarding any excess bits beyond that precision) will provide the correct result. To illustrate multiplication we consider the following example: $6 \times -5 = -30$ and $n=4$. First, the input operands are sign-extended from 4 bits to 8. Then the numbers are multiplied, discarding the bits beyond bit at position 7 (shown by 'x'):

Example 4-8. Find $A * B$ if $A=6$ and $B=-5$, $n=4$.

	0000 0110	(6)	Multiplicand
	× 1111 1011	(-5)	Multiplier ($M_7M_6M_5M_4M_3M_2M_1M_0$)
=====			
	110		Multiplicand* M_0 * 2^0
+	110		Multiplicand* M_1 * 2^1
+	0		Multiplicand* M_2 * 2^2
+	110		Multiplicand* M_3 * 2^3
+	110		Multiplicand* M_4 * 2^4
+	110		Multiplicand* M_5 * 2^5
+	x10		Multiplicand* M_6 * 2^6
+	xx0		Multiplicand* M_7 * 2^7
=====			
P:	x11110010	(-30)	

4.6 Binary Coded Decimal Numbers

Binary-coded decimal (BCD) is an encoding for decimal numbers in which each decimal digit is represented by its own binary sequence. Its main virtue is that it allows easy conversion to decimal digits for printing or display and faster decimal calculations. Its drawbacks are the increased complexity of the logic circuits needed to implement mathematical operations and a relatively inefficient encoding—it occupies more space than a pure binary representation. For example, a byte can represent positive integers in range from 0 to 99 if using packed BCD representation, whereas regular binary represents positive integers in range from 0 to 255. Though BCD is not as widely used as it once was, decimal fixed-point and floating-point are still important and still used in financial, commercial, and industrial computing.

To BCD-encode a decimal number using the common encoding, each decimal digit is stored in a four-bit nibble.

Decimal:	0	1	2	3	4	5	6	7	8	9
BCD:	0000	0001	0010	0011	0100	0101	0110	0111	1000	1001

Thus, the BCD encoding for the number 127 would be: 0001_0010_0111

Since most computers store data in eight-bit bytes, there are two common ways of storing four-bit BCD digits in those bytes:

- each digit is stored in one nibble of a byte, with the other nibble being set to all zeros, all ones (as in the EBCDIC code), or to 0011 (as in the ASCII code)
- two digits are stored in each byte (so called packed BCD encoding).

Unlike binary encoded numbers, BCD encoded numbers can easily be displayed by mapping each of the nibbles to a different character. Converting a binary encoded number to decimal for display is much harder involving integer multiplication or divide operations.

5 Characters

Computers represent not only numbers in binary, but also characters, whether they are letters, digits, or special characters. All characters that a computer can use are called a character set. Two common character sets are ASCII – American Standard Code for Information Interchange and Unicode. In this course we will exclusively focus on ASCII.

ASCII was developed in the 1960s and was based on earlier codes used by telegraph systems. ASCII code defines 128 characters, which are assigned numbers from 0 to 127 using seven bits. Characters are represented in a table as shown in Table 3.

We use one byte to store each character (bit 7 in a byte is not used). You will notice that ASCII code includes non-printable characters that control how text is processed (0 – 31 and 127) and printable characters (32-126). Among printable characters you will find special characters (e.g., ASCII code for [SPACE] is 32 or 0x20), digits (e.g., ASCII code for character '0' is 48 or 0x30, '1' is 49 or 0x31, etc), upper case letters (ASCII code for character 'A' is 65 in decimal or 0x41 in hexadecimal), and lower case letters (ASCII code for character 'a' is 97 or 0x61).

Let us consider a string or array of characters "CPE 323". Please note that this string has 7 characters, 'C', 'P', 'E', ' ', '3', '2', '3'. In memory it will occupy 7 bytes: the first byte represents upper case letter 'C' and is encoded as 0x43, the second byte represents upper case letter 'P' and is encoded as 0x50, the third byte represents 'E' and is encoded as 0x45, the fourth byte encodes the space character and is encoded as 0x20, the fifth byte represents '3' and is encoded as 0x33, the sixth byte is '2' encoded as 0x32, and the last byte is '3' encoded as 0x33.

Check yourself 5-1.

How many bytes is used to encode a message "Hello" in ASCII and what is their content in hexadecimal number system?

Table 3. ASCII Table.

Decimal	Hex	Char	Decimal	Hex	Char	Decimal	Hex	Char	Decimal	Hex	Char
0	0	[NULL]	32	20	[SPACE]	64	40	@	96	60	`
1	1	[START OF HEADING]	33	21	!	65	41	A	97	61	a
2	2	[START OF TEXT]	34	22	"	66	42	B	98	62	b
3	3	[END OF TEXT]	35	23	#	67	43	C	99	63	c
4	4	[END OF TRANSMISSION]	36	24	\$	68	44	D	100	64	d
5	5	[ENQUIRY]	37	25	%	69	45	E	101	65	e
6	6	[ACKNOWLEDGE]	38	26	&	70	46	F	102	66	f
7	7	[BELL]	39	27	'	71	47	G	103	67	g
8	8	[BACKSPACE]	40	28	(72	48	H	104	68	h
9	9	[HORIZONTAL TAB]	41	29)	73	49	I	105	69	i
10	A	[LINE FEED]	42	2A	*	74	4A	J	106	6A	j
11	B	[VERTICAL TAB]	43	2B	+	75	4B	K	107	6B	k
12	C	[FORM FEED]	44	2C	,	76	4C	L	108	6C	l
13	D	[CARRIAGE RETURN]	45	2D	-	77	4D	M	109	6D	m
14	E	[SHIFT OUT]	46	2E	.	78	4E	N	110	6E	n
15	F	[SHIFT IN]	47	2F	/	79	4F	O	111	6F	o
16	10	[DATA LINK ESCAPE]	48	30	0	80	50	P	112	70	p
17	11	[DEVICE CONTROL 1]	49	31	1	81	51	Q	113	71	q
18	12	[DEVICE CONTROL 2]	50	32	2	82	52	R	114	72	r
19	13	[DEVICE CONTROL 3]	51	33	3	83	53	S	115	73	s
20	14	[DEVICE CONTROL 4]	52	34	4	84	54	T	116	74	t
21	15	[NEGATIVE ACKNOWLEDGE]	53	35	5	85	55	U	117	75	u
22	16	[SYNCHRONOUS IDLE]	54	36	6	86	56	V	118	76	v
23	17	[ENG OF TRANS. BLOCK]	55	37	7	87	57	W	119	77	w
24	18	[CANCEL]	56	38	8	88	58	X	120	78	x
25	19	[END OF MEDIUM]	57	39	9	89	59	Y	121	79	y
26	1A	[SUBSTITUTE]	58	3A	:	90	5A	Z	122	7A	z
27	1B	[ESCAPE]	59	3B	;	91	5B	[123	7B	{
28	1C	[FILE SEPARATOR]	60	3C	<	92	5C	\	124	7C	
29	1D	[GROUP SEPARATOR]	61	3D	=	93	5D]	125	7D	}
30	1E	[RECORD SEPARATOR]	62	3E	>	94	5E	^	126	7E	~
31	1F	[UNIT SEPARATOR]	63	3F	?	95	5F	_	127	7F	[DEL]

6 Fraction Numbers

To represent real numbers in computers two major approaches are used: fixed-point and floating-point representation.

6.1 Fixed-point

For fixed-point representation, the binary radix point is assigned to a fixed location in a byte (or word). Fixed-point numbers are useful for representing fractional values when the executing processor has no floating-point unit (FPU) and when fixed-point representation provides improved performance and sufficient accuracy for the application at hand. Most low-cost embedded microprocessors and microcontrollers do not have an FPU.

In binary fixed-point numbers, each magnitude bit represents a power of two, while each fractional bit represents an inverse power of two. Thus the first fractional bit is $\frac{1}{2}$, the second is $\frac{1}{4}$, the third is $\frac{1}{8}$ and so on. Signed fixed-point numbers in two's complement format are defined as follows, where m and f are the number of bits in the integer portion (M) and the fraction portion (F), respectively:

$$A = A_{m-1}A_{m-2} \dots A_0.A_{-1}A_{-2} \dots A_{-f} = -A_{m-1} \cdot 2^{m-1} + A_{m-2} \cdot 2^{m-2} + \dots + A_0 \cdot 2^0 + A_{-1} \cdot 2^{-1} + \dots + A_{-f} \cdot 2^{-f}$$

The upper bound for this representation is $2^{(m-1)}-2^{(-f)}$ and the lower bound is given by $-2^{(m-1)}$.

Unsigned fixed-point numbers are defined as follows (the range is from 0 to $2^m - 2^{-f}$):

$$A = A_{m-1}A_{m-2} \dots A_0.A_{-1}A_{-2} \dots A_{-f} = A_{m-1} \cdot 2^{m-1} + A_{m-2} \cdot 2^{m-2} + \dots + A_0 \cdot 2^0 + A_{-1} \cdot 2^{-1} + \dots + A_{-f} \cdot 2^{-f}$$

Let's consider the following unsigned number with 5-bit integer portion and 3-bit fraction:

$$00000.101_2 = 0.625_{10}$$

Granularity of precision is a function of the number of fractional bits assigned. Thus, if $f=3$, the smallest fraction is $2^{-3} = 0.125$. The maximum number that can be represented in this notation (8 bits, 5 for integer portion, 3 for fraction) is 63.875 (encoded as 11111.111) and the minimum is 0.0 (encoded as 00000.000).

6.2 Floating-point

The IEEE Standard for Binary Floating-Point Arithmetic (IEEE 754) is the most widely-used standard for floating-point computation, and is followed by many CPU (Central Processing Unit) and FPU (Floating-Point Unit) implementations. The standard defines formats for representing floating-point numbers (including negative zero and denormal numbers) and special values (infinities and NaNs), together with a set of floating-point operations that operate on these values. It also specifies four rounding modes and five exceptions (including when the exceptions occur, and what happens when they do occur).

Floating-point numbers are typically packed into a computer datum as the sign bit, the exponent field, and the significand (mantissa), from left to right. For the IEEE 754 binary formats they are apportioned as shown in Table 4.

Table 4. Floating-point types as defined by the IEEE 754 standard.

Type	Sign	Exponent	Exponent bias	Significand	Total
Half (IEEE 754-2008)	1	5	15	10	16
Single	1	8	127	23	32
Double	1	11	1023	52	64
Quad	1	15	16383	112	128

Let's take a look at a single precision floating-point format that requires 32 bits as follows: $S_EEEEEEEE_FFFFFFFFFFFFFFFF$, where S is the sign bit, Es represent the exponent bits, and Fs represent the mantissa bits:

S: sign bit, E: exponent bits; F: mantissa (significand) bits

position 31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0
field S E₇ E₆ E₅ E₄ E₃ E₂ E₁ E₀ F₋₁ F₋₂ F₋₃ F₋₄ F₋₅ F₋₆ F₋₇ F₋₈ F₋₉ F₋₁₀ F₋₁₁ F₋₁₂ F₋₁₃ F₋₁₄ F₋₁₅ F₋₁₆ F₋₁₇ F₋₁₈ F₋₁₉ F₋₂₀ F₋₂₁ F₋₂₂ F₋₂₃

The 32-bit representation consists of three parts. The most significant bit indicates the sign of the number. The next 8 bits are reserved for the exponent, and the last 23 bits are used for the fraction.

The value of the single-precision number is determined as follows:

$$\text{Value} = (-1)^S 2^{(E-127)} * (1.F)$$

Converting decimal digits to IEEE binary floating point is a little tricky. There are 3 steps.

- The first step in the conversion is the simplest. This is determining the sign bit. If the number is positive, then the sign bit is 0, $S=0$, otherwise it is 1, $S=1$.
- The next eight digits are used to express the exponent, which we will figure out last.
- The final 23 digits are used to express the fraction.

Let us first illustrate the method for the conversion using a simple example. We will use -210.25 to walk through the conversion. Determining the sign bit is easy – this is a negative number, so $S=1$.

The next step is to convert the absolute value of your decimal number to binary (210.25). It is easiest to focus on the integer portion of the number first, then the fraction. 210 is 11010010 or $128+64+16+2$, otherwise expressed as:

$$210_{10} = 1 \cdot 2^7 + 1 \cdot 2^6 + 0 \cdot 2^5 + 1 \cdot 2^4 + 0 \cdot 2^3 + 0 \cdot 2^2 + 1 \cdot 2^1 + 0 \cdot 2^0 = 11010010_2$$

Next, we need to convert the fractional part. To do this we have to convert the number into a binary sequence that takes the following form:

$$A_{-1} \cdot 2^{-1} + A_{-2} \cdot 2^{-2} + A_{-3} \cdot 2^{-3} + \dots$$

Luckily .25 is $1/4$ and so this is easy to convert it to binary:

$$0.25_{10} = 0 \cdot 2^{-1} + 1 \cdot 2^{-2} = .01_2$$

Thus, the decimal 210.25 is represented as 11010010.01 in binary.

$$210.25_{10} = 11010010.01_2 = 1.101001001 \cdot 2^7$$

Note: generally, converting the fractional part may include multiple steps as shown in the example below.

Example 6-1. Convert the fractional part 0.1 into binary.

Converting the fractional part involves multiplying by 2 in each step. The integer portion of the product is kept as a new digit in the binary representation and the fractional part of the product is used in the next step. The process of multiplying by 2 continues until you get the fractional part of the product equal to 0. Some fractions cannot be represented finitely in the binary system. In such cases the process stops when the desired number of bits is found. The fraction 0.1 is such an example that cannot be represented in binary with a finite number of bits.

Operation	Bits	(Position)
$0.1 * 2 = 0.2$	0	(bit A ₋₁ : weight 2^{-1})
$0.2 * 2 = 0.4$	0	(bit A ₋₂)
$0.4 * 2 = 0.8$	0	(bit A ₋₃)
$0.8 * 2 = 1.6$	1	(bit A ₋₄)
$0.6 * 2 = 1.2$	1	(bit A ₋₅)
$0.2 * 2 = 0.4$	0	(bit A ₋₆)
$0.4 * 2 = 0.8$	0	(bit A ₋₇)
$0.8 * 2 = 1.6$	1	(bit A ₋₈)
. . .		
$0.1_{10} = 0.00011001\dots_2 = 1*2^{-4} + 1*2^{-5} + 1*2^{-8} + 1*2^{-9}$		

The next step is to normalize this number so that only one non-zero decimal place is in the non-fraction part of the number. To do this we must shift the decimal place 7 positions to the left (see above). The number 7 becomes important so we note it. This process leaves us with the number 1.101001001, which is the normalized form 1.F. The fraction F is represented in the last 23 bit places in the 32 bit binary, F=101001001. The fraction is padded with 0's to the right to fill in the full 23 bits - leaving us with F=10100100100000000000000 for the mantissa or significand.

So we now have the first bit S, S=1, and the last 23 bits of the 32 bit sequence. We must now derive the middle 8 bits. To do this we take our exponent (7) and add 127 (the exponent bias for single-precision numbers) to get E=134. We then express this number as an 8 bit binary. This is 10000110 (or 128+4+2). Now we have the middle bits and can stitch together all bits to get a complete single-precision binary representation:

$$1_10000110_10100100100000000000000_2 = C3524000_{16}$$

We can convert this bit sequence back into a number by reversing the process. First we can note by the leading 1 that the number is negative. Next we can determine the exponent. 10000110 is binary for 2+4+128 or 134. 134-127 is 7, so the exponent is 7. Finally we take the last 23 digits, convert them back into the original fraction (adding the preceding 1.) to get:

$$1.101001001$$

Moving the decimal place to the right by 7 (corresponding to the exponent) we get:

11010010.01

This binary is equal to $128+64+16+2 + 1/4$ or 210.25. Once we apply the negative sign (indicated by the leading bit set to 1) we get our original number:

-210.25

Special Values

The IEEE-754 standard reserves the exponent field values $E=0$ (0000_0000) and $E=255$ (1111_1111) to denote special values in the floating-point representation.

Zero

As mentioned above, zero is not directly representable in the straight format, due to the assumption of a leading 1 in the normalized form (we need to specify a true zero mantissa to yield a value of zero). Thus, zero is a special value denoted with an exponent field of zero and a fraction field of zero. Note that -0 and $+0$ are distinct values, though they both compare as equal.

Denormalized Numbers

If the exponent $E=0$ (0000_0000), but the fraction is non-zero, then the value is a denormalized number, which does not have an assumed leading 1 before the binary point. Thus, the value of the numbers represented is determined as follows: $(-1)^s \times 0.F \times 2^{-126}$. For double precision, denormalized numbers their value is determined as follows: $(-1)^s \times 0.F \times 2^{-1022}$. From this you can interpret zero as a special type of a denormalized number.

Infinity

The values $+\infty$ and $-\infty$ are denoted with an exponent of all 1s and a fraction of all 0s ($E=255, F=0$). The sign bit distinguishes between negative infinity and positive infinity. Being able to denote infinity as a specific value is useful because it allows operations to continue past overflow situations. Operations with infinite values are well defined in the IEEE floating-point standard.

Not A Number

The value NaN (Not a Number) is used to represent a value that does not represent a real number. NaN's are represented by a bit pattern with the exponent $E=255$ (1111_1111) and a non-zero fraction. There are two categories of NaN: QNaN (Quiet NaN) and SNaN (Signaling NaN).

A QNaN is a NaN with the most significant fraction bit set. QNaN's propagate freely through most arithmetic operations. These values pop out of an operation when the result is not mathematically defined.

An SNaN is a NaN with the most significant fraction bit cleared. It is used to signal an exception when used in operations. SNaN's can be handy to assign to uninitialized variables to trap premature usage.

Semantically, QNaNs denote indeterminate operations, while SNaNs denote invalid operations.

Summary

Table 5 summarizes the IEEE 754 floating-point representation for fractional numbers.

Table 5. Floating-point representation (b is bias).

Sign (s)	Exponent (e)	Fraction (f)	Value
0	00 ... 00	000...0	+0
0	00 ... 00	00 ... 01 11 ... 11	Positive denormalized real $0.f \times 2^{-(b+1)}$
0	00 ... 01 11 ... 10	xx ... xx	Positive normalized real $1.f \times 2^{(e-b)}$
0	11 ... 11	00 ... 00	+Infinity
0	11 ... 11	00 ... 01 01 ... 11	SNaN
0	11 ... 11	10 ... 00 11 ... 11	QNaN
1	00 ... 00	000...0	-0
1	00 ... 00	00 ... 01 11 ... 11	Negative denormalized real $-0.f \times 2^{-(b+1)}$
1	00 ... 01 11 ... 10	xx ... xx	Negative normalized real $-1.f \times 2^{(e-b)}$
1	11 ... 11	00 ... 00	-Infinity
1	11 ... 11	00 ... 01 01 ... 11	SNaN
1	11 ... 11	10 ... 00 11 ... 11	QNaN

7 References

- http://en.wikipedia.org/wiki/Binary_number
- http://en.wikipedia.org/wiki/Two's_complement
- http://en.wikipedia.org/wiki/IEEE_754
- <https://en.wikipedia.org/wiki/ASCII>
- Web IEEE-754 Floating-point Converter:
<https://www.h-schmidt.net/FloatConverter/IEEE754.html>

8 Exercises

Problem #1.

Fill in the following table. Assume 16-bit memory locations. You should provide binary, hexadecimal, decimal (as positive integer), and binary coded decimal representations. Show your work as illustrated for (i). Note: for 4-bit BCD representation, if the digit value is outside the range 0-9 use '?'.

	16-bit binary	Hexadecimal (4 hex digits)	Decimal (unsigned int)	2-byte packed BCD (4-BCD digits)
(i)	0011.0000.0011.1100	303C	<u>12,348</u>	303?
(a)		<u>cbf3</u>		
(b)			9,080	
(c)	<u>1001.0100.0010.1100</u>			
(d)		<u>1943</u>		

(i)

Here 12,348 decimal representation is provided. We convert this number into a hexadecimal number by dividing by 16 as follows.

$$12348/16 = 771 \quad \lfloor 12$$

$$771/16 = 48 \quad \lfloor 3$$

$$48/16 = 3 \quad \lfloor 0$$

$$3/16 = 0 \quad \lfloor 3$$

From hexadecimal representation we can get directly binary and BCD representations.

$$12348_{10} = 303C_{16} = 0011_0000_0011_1100_2 = 303? \text{ ("?" marks an illegal BCD digit).}$$

Problem #2.

What is the range of unsigned and signed integers that can be represented using 16 bits?

Range of 16-bit unsigned integers: _____

Range of 16-bit signed integers (in 2's complement): _____

Consider the following 16-bit hexadecimal numbers given in the second column of the table below. Each of these values can be interpreted as an unsigned 16-bit integer or a signed 16-bit integer represented in 2's complement.

Provide the decimal value for each number and interpretation. Show your work as illustrated in (i).

	16-bit hex	Unsigned int	Signed int (2's complement)
(i)	A223	41507	-24029
(a)	81C2		
(b)	B607		
(c)	39CD		
(d)	0012		

(i) unsigned: $A223_{16} = 10 \cdot 16^3 + 2 \cdot 16^2 + 2 \cdot 16^1 + 3 \cdot 16^0 = 41,507_{10}$

signed: $A223_{16} = 1010.0010.0010.0011_2 \Rightarrow$ this is a negative number;

two's complement is: $0101.1101.1101.1101 = 5DDD_{16} = 24,029_{10} \Rightarrow A223_{16} = -24,029$

Problem #3.

Consider the following arithmetic operations. Find the results and set the flags C, V, N, and Z accordingly. Show your work using hexadecimal or binary representation.

(a) 8-bit, two's complement

$$24_{10} + 105_{10}$$

(b) 8-bit, two's complement

$$(-55)_{10} - 68_{10}$$

(c) 16-bit, two's complement

$$(-45)_{10} + 88_{10}$$

Problem #4.

(a) Convert the following number from decimal to a single-precision IEEE-754 floating point number.

$$-32.125_{10}$$

(b) Convert the following IEEE-754 32-bit floating point number from hex/binary to decimal.

$$60E3AB00_{16}$$

(c) What data type would you need to use to represent the following number, 1,048,576.0625, without losing any significant bits of the mantissa? Explain your answer.

(d) What is the smallest negative normalized half-precision floating-point number? Give its decimal representation. Show your work.

(e) What is the largest positive normalized half-precision floating-point number? Give its decimal representation. Show your work.

Problem #5.

(a) Convert the following number from decimal to a single-precision IEEE-754 floating point number. Can you accurately represent this number in single-precision?

0.1_{10}

(b) What is a true value in your single-precision representation? What is the error due to conversion?

(c) Will double-precision IEEE-754 provide an accurate representation of 0.1? If no, what is the error in that case?

