

Fundamentals of Programming 1

Enumerations and Linear Arrays

Arkadiusz Chrobot

Department of Information Systems

November 29, 2024

Outline

- 1 Data Abstraction
- 2 Enumerations
- 3 The `typedef` Keyword
- 4 Linear Arrays
- 5 Initialization of an Array
- 6 Basic Operations on Linear Array

Data Abstraction

Abstraction can be applied not only to statements but also to data. The C programming language allows a programmer to create custom data types specific to a problem that is to be solved. An example of such a data type is the enumerated type.

Enumerated Types

An enumeration makes it possible to name characters or numbers that belong to a subset of integers. In other words, the enumeration may be considered as a set of constants. It has to have its data type. The overall pattern for defining an enumerated type is as follows:

```
enum type_name {ELEMENT_1=value, ..., ELEMENT_N};
```

Please note, that the name (identifier) of each element is written in uppercases, just like names of constants. This is only a matter of convention. Any name can be used, provided it is valid in the C language. To each element of the enumerated type can be assigned a character or an integer number. However, if no value is given to the elements by the programmer, the compiler will assign to them successive numbers of `int` type, starting from 0. It is also possible for the programmer to assign the same value to more than one element, or to assign a value only to one element or a selected group or groups of elements. The rest of them will get default values.

Enumerated Types

Examples

```
enum names_of_days {MONDAY=0, TUESDAY=1, WEDNESDAY=2, THURSDAY=3,  
                    FRIDAY=4, SATURDAY=5, SUNDAY=6};
```

In the above definition of an enumerated type an integer number is assigned to each name of a day. The same definition can be written as follows:

```
enum names_of_days {MONDAY, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY,  
                    SATURDAY, SUNDAY};
```

However, if a programmer wants the values of the days to start with 1 instead of 0 she or he should define the type like this:

```
enum names_of_days {MONDAY=1, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY,  
                    SATURDAY, SUNDAY};
```

Each subsequent day, after Monday, gets a number greater by one than its predecessor, i.e. TUESDAY=2, WEDNESDAY=3, etc.

Enumerated Types

Examples

If the programmer wants to distinguish the days of weekend, by assigning them for example numbers 9 and 10, then she or he can do that like this:

```
enum names_of_days {MONDAY=1, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY,  
                    SATURDAY=9, SUNDAY};
```

The values can also be assigned to the elements in a descending order:

```
enum directions {NORTH=3, WEST=2, EAST=1, SOUTH=0};
```

Enumerations

An enumeration (a variable of an enumerated type) may be declared as local (also as a function's parameter) or global. The overall pattern for such a declaration is as follows:

```
enum name_of_enumerated_type name_of_variable;
```

Any element of its enumerated type can be assigned to the variable. The enumeration may be applied as a loop counter in a `for` loop, a selector in a `switch` statement, or it may be used in conditional statements and in the condition-controlled loops.

Unfortunately, the C language implementation of enumerations is imperfect. They are only a facilitation for the programmer. The compiler does not verify the correctness of using the enumerated type variables, considering them to be of `int` type. Thus, it is possible to assign to an enumeration any integer number. This can lead to many mistakes. The C language makes it also possible to use the elements of enumerated type as common constants and to define constants of enumerated type with the use of the `const` keyword.

Enumerations

Examples

```
#include <stdio.h>

enum names_of_days {MONDAY, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY, SATURDAY, SUNDAY};

void print_message(enum names_of_days day)
{
    switch(day) {
        case MONDAY:
        case TUESDAY:
        case WEDNESDAY:
        case THURSDAY:
        case FRIDAY:
            puts("Go to work!");
            break;
        default:
            puts("Relax!");
    }
}

int main(void)
{
    enum names_of_days day;
    for(day=MONDAY; day<=SUNDAY; day++)
        print_message(day);
    return 0;
}
```


Enumerations

Comment to the Example

In the program from the previous slide, a local enumerated type variable is declared in the `main()` function. It is the `day` variable. Also in the `print_message()` function a parameter of the same type and name is declared. The variable in the `main()` function is used as a loop counter and the one in the `print_message()` is applied as a `switch` selector. It is worth noticing, that in the `switch` statement most of the cases are empty, but it is not a mistake but an intended effect. That way, the code for Friday is performed also for all other days, except for Saturday and Sunday, and it is not repeated. The code for those two mentioned days is performed in the default case.

There is no easy way to print an element of an enumeration. It is however possible to display its numerical value with the use of `printf()` function and the `%d` conversion specifier.

Enumerations

Examples

```
#include <stdio.h>

enum names_of_days {MONDAY, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY, SATURDAY=9, SUNDAY};

void print_message(enum names_of_days day)
{
    switch(day) {
        case MONDAY:
        case TUESDAY:
        case WEDNESDAY:
        case THURSDAY:
        case FRIDAY:
            puts("Go to work!");
            break;
        default:
            puts("Relax!");
    }
}

int main(void)
{
    enum names_of_days day;
    for(day=MONDAY; day<=SUNDAY; day++)
        print_message(day);
    return 0;
}
```

Enumerations

Comment to the Example

A small change in the program, that assigns the value of 9 to the `SATURDAY` element shows the imperfections of enumerations. The output of the program suggests that the week has 11 days, most of them free. The problem here is the `day` variable which is applied as a loop counter. When the value of the counter is 4, what corresponds to the `FRIDAY` element, it is incremented in the next iteration of the loop to 5 which cannot be mapped to any element of the enumerated type, but it is still a proper value, because for the program the `day` variable is of the `int` type. In the C language the enumeration is only a simple container for constants and should be used carefully.

The typedef Keyword

It is easy to forget that the declaration of an enumeration always starts with the `enum` keyword. To avoid such an issue the `typedef` keyword may be applied. It allows for giving the enumerated type a shorter name:

```
typedef enum names_of_days {MONDAY, TUESDAY, WEDNESDAY, THURSDAY, FRIDAY, SATURDAY=9, SUNDAY} days;
```

The variable of the type may now be declared like this:

```
days day = MONDAY;
```

Generally, the `typedef` keyword makes it possible to give an alternative name to any data type including the predefined data types of the C language. Therefore, it should be used carefully. Many programmers advise against using it, since it can make the program difficult to read by concealing the true types of variables.

Linear Arrays

One-dimensional arrays or simply the linear arrays are an example of a *data structure* i.e. a variable that can store more than one value at the same time. In case of arrays all these values are of the same type. The picture below shows the construction of a linear array that can store up to 8 integer numbers.

0	1	2	3	4	5	6	7
-7	7	10	-3	0	7	15	0

Arrays are build from *elements*. Numbers on the top of the array are the indices that uniquely locate an element inside the array. In the C language the indices are natural numbers. The first element of the array always has an index of 0. The element is a single location in the array that holds a single value. All indices are unique and form an ascending sequence. The values of the elements may repeat and do not have to be sorted. The array is sometimes called a *vector*.

Linear Arrays

Declaration

Just like other variables the array has to be declared before using. It can be declared as global or local variable. The elements of the former are initialized with zeros and the initial value of the elements of the latter is unspecified. The overall pattern of declaration of an array is as follows:

```
data_type array_name [NUMBER_OF_ELEMENTS] ;
```

An element of the array may be of any of the types that were already introduced in the lecture. The arrays of characters have a special meaning and they will be discussed separately in a future lecture. The name (identifier) has to be legal in the C language. The number of elements, also called the *length of an array*, determines how many elements will have the array. It is given as a *constant expression*, like a literal or a constant defined with the use of `#define` statement. According to the ISO C99 standard, the number of elements has to be greater than zero. The indices range starts from 0 and ends with `NUMBER_OF_ELEMENTS-1`.

Linear Arrays

Accessing the elements

It is worth noticing, that the array resembles the structure of RAM. In order to read or modify the content of a single memory cell an address has to be provided. Similarly, if a programmer wants to read or write an element of an array she or he has to use its index. Referencing an element of an array follows this pattern:

```
array_name[index]
```

In the C language the array name is equivalent to a pointer. Thus there are three additional ways of referencing the element:

```
*(array_name+index)
```

```
*(index+array_name)
```

```
index[array_name]
```

The first two expressions use so-called pointers arithmetic. The most frequently used way of accessing an array element is the one introduced as first. The second one is also sometimes applied. The last two are rarely used because they are less legible.

Linear Array

Array Size and Number of Elements

The size of an element (the number of bytes it occupies in memory) can be acquired with the use of `sizeof` operator. The number of array's elements can be computed using the following expression:

```
sizeof(array_name)/sizeof(array_name[0])
```

The expression can be a little simplified with the use of the pointer arithmetic:

```
sizeof(array_name)/sizeof(*array_name)
```

Unfortunately, both expressions and the `sizeof` operator give incorrect results when the array is a function's parameter. In this case it returns the size of array's single element. Fixing such an issue is possible and it will be described latter.

Arrays And Functions

Passing to a Function

Functions make implementing operations on arrays easy. For the rest of the lecture only single functions operating on arrays are presented, instead of entire programs. Arrays can and should be passed by parameters to functions. The array parameter can be declared in the same way as a regular array, although in a function's list of parameters. An example of a program that uses such a parameter is given in the next slide. The array is passed as an argument to the `print()` function, when it is called in the `main()` function. To this end the array's name is put in the parentheses in the invocation of `print()`. This function displays the size of the array parameter. It is equal to the size of pointer, and not the array. The compiler informs about that by issuing a warning.

Arrays And Functions

Passing to a Function

```
#include<stdio.h>
#define NUMBER_OF_ELEMENTS 10

int array[NUMBER_OF_ELEMENTS];

void print(int array_parameter[NUMBER_OF_ELEMENTS])
{
    printf("Size: %lu\n", sizeof(array_parameter));
}

int main(void)
{
    print(array);
    return 0;
}
```

Arrays And Functions

Passing to a Function

Number of elements in the array parameter is ignored. It means that an array of any number of elements can be passed by a parameter that, for example, has 10 elements specified in its definition, as long as the data types of the array's and parameter's elements are compatible. Therefore, the number of elements in the definition of array parameter is usually omitted — the brackets are left empty. In the next slide is included the same program as before, but without the number of elements in the parameter.

Arrays And Functions

Passing to a Function

```
#include<stdio.h>
#define NUMBER_OF_ELEMENTS 10

int array[NUMBER_OF_ELEMENTS];

void print(int array_parameter[])
{
    printf("Size: %lu\n", sizeof(array_parameter));
}

int main(void)
{
    print(array);
    return 0;
}
```

Arrays And Functions

Passing to a Function

An array may also be passed to the function with the use of a pointer parameter. The data type of the pointer should be the same as the data type of the elements of the array, or it should be `void *`. The second data type means that the pointer can be assigned an address of any variable, including an array. The next slide contains the modified program from the previous slide, where the array is passed by a pointer. **Effectively, any array is always passed by a pointer, regardless of what kind of parameter is used for that operation. It means that changing a value of the parameter's element results in the modification of the value of the same element in the array passed as argument.**

Arrays And Functions

Passing to a Function

```
#include<stdio.h>
#define NUMBER_OF_ELEMENTS 10

int array[NUMBER_OF_ELEMENTS];

void print(int *array_parameter)
{
    printf("Size: %lu\n", sizeof(array_parameter));
}

int main(void)
{
    print(array);
    return 0;
}
```

Arrays And Functions

Passing to a Function

To determine, using the formulas introduced earlier, how many elements has an array passed to the function, we need to define an array type with the use of the `typedef` keyword and then declare a parameter of this type. In that case, we can use the `sizeof` operator inside the function, to get the size of the array. The argument of this operator should be the array type. This solution is applied in the program presented in the next slide.

Arrays And Functions

Passing to a Function

```
#include<stdio.h>
typedef int int_array_type[10];
int_array_type array;

void print(int_array_type array_parameter)
{
    printf("%lu\n",
           sizeof(int_array_type)/sizeof(*array_parameter));
}

int main(void)
{
    print(array);
    return 0;
}
```


Arrays And Functions

Returning Array From Function

Arrays cannot be returned directly from a function, but their addresses can, provided that the data type of the function's returned value is declared as an appropriate pointer. **Addresses of local arrays, or any other local variables must never be returned from a function, because as soon as the function completes, the variables cease to exist.** The exception to this rule are local variables declared with the use of the `static` keyword.

Arrays And Functions

Variable-Length Array

The ISO C99 standard has introduced *variable-length arrays* (VLAs) to the language. They are local arrays whose length is determined at a runtime. Although that sounds like a good idea, it can also be a source of many issues. If the number of the elements is too large, then the size of the array will exceed the size of the call stack and that will result in a program crash. The VLAs cannot be initialized at the place of their declaration, because their number of elements cannot be determined at the time of compilation. For that reasons, VLAs have been made optional in the ISO C11 standard. A program that creates such an array is included in the next slide. Since it doesn't use this array, the compiler issues a warning when it is compiled.

Arrays And Functions

Variable-Length Arrays

```
void create_vla(unsigned int number_of_elements)
{
    if(number_of_elements>0 && number_of_elements<1000) {
        int array[number_of_elements];
    }
}

int main(void)
{
    create_vla(100);
    return 0;
}
```

Initialization of an Array

Initialized Array

The array may be initialized in the place of its declaration. To this end an assignment operator has to be put after the closing bracket followed by a list of values in curly braces. The values have to be separated by commas. If there is provided less values in the list than the array has elements then the last elements of the array will be initialized with zeros. It is also possible to skip the number of elements when initializing an array this way. The compiler will figure it out by the number of provided initial values.

```
int main(void)
{
    double fractions[] = {0.1, 0.2, 0.3};
    double fractions_2[3] = {0.1, 0.2, 0.3};
    return 0;
}
```

Initialized of an Array

Designated Initializers

The simplest way of initializing all elements of a local array with zeros is demonstrated in the following code:

```
int main(void)
{
    int array[100]={0};
    return 0;
}
```

Starting with the ISO C99 standard the C language offers *designated initializers* that make it possible to specify an element, or group of elements, to which a given value has to be assigned.

```
int main(void)
{
    int array[100]={ [5]=2, [3]=1};
    return 0;
}
```

Initialization of an Array

Initialization Made by User

The initial values for the array can be provided by user, during a program run, with the use of a keyboard. It requires using a loop for indexing the elements of an array and the `scanf()` function. An element of the array is a single variable. Therefore its address should be provided as the second argument of the `scanf()` call. The address is obtained with the use of the address operator. An example:

```
int main(void)
{
    int array[5];
    unsigned int i;

    for(i=0;i<5;i++)
        scanf("%d",&array[i]);

    return 0;
}
```

Initialization of an Array

Initialization with the Use of Indices

Using one of the two introduced ways of initializing arrays in case where the array has huge number of elements could be cumbersome. Alternatively, if the elements of the array are of `int` or compatible type, they can be assigned the values of their indices:

```
int main(void)
{
    int array[1000];
    unsigned int i;
    for(i=0;i<1000;i++)
        array[i] = i;
    return 0;
}
```

It is a simple way of initializing an array. However, the values of the elements are unique and arranged in an ascending sequence, which is not always desired.

Pseudorandom Number Generator

An array may be initialized with the use of a pseudorandom number generator (PRNG). It is an algorithm that generates numbers appearing to be random. However, it can be statistically verified that they are not truly random. This usually excludes their use in cryptography, but for the purpose of this lecture they are sufficiently random. To distinguish them from truly random numbers, they are called pseudorandom. The algorithm uses an initial value called a *seed* and a mathematical formula to generate such numbers.

Pseudorandom Number Generator

Using the PRNG in the C Language

In the C language there are two functions, declared in the `stdlib.h` header file, that are the implementation of the PRNG. These are `srand()` and `rand()`. The former takes as an argument an `unsigned int` number and makes it the seed for the PRNG. The latter does not take any argument and returns a pseudorandom `int` number ranging from 0 to `RAND_MAX`. The `srand()` function **has to be called outside any loop**. The argument for the function may be the result of the `time()` function declared in the `time.h` header file. An argument for the latter function can be `NULL` or zero.

Pseudorandom Number Generator

The PRNG Usage

The PRNG generates only natural numbers. If a pseudorandom natural number is needed ranging from 0 to 9 (both inclusive) it can be chosen with the following expression:

```
int x = rand()%10;
```

To choose numbers ranging from 1 to 10 (both inclusive) the expression should be modified as follows:

```
int x = 1+rand()%10;
```

To pick integer numbers ranging from -10 to 10 (both inclusive), a following expression can be applied:

```
int x = -10+rand()%21;
```

If a floating-point pseudorandom number is needed from the interval of $[-10,11)$ it could be chosen with the following expression:

```
double x = -10+rand()%21+rand()/(RAND_MAX+1.0);
```

Pseudorandom Number Generator

The PRNG Usage

To pick a pseudorandom lowercase letter from the set of 26 lowercases the following expression may be used:

```
char x = 'a'+rand()%26;
```

Array Initialization

Array Initialization with Pseudorandom Numbers

Below is defined a function that initializes an array of `NUMBER_OF_ELEMENTS` elements with pseudorandom numbers ranging from 0 to 199 (both inclusive):

```
void fill_array_with_random_numbers(int array[])
{
    srand(time(0));
    int i;
    for(i=0; i<NUMBER_OF_ELEMENTS; i++)
        array[i]=rand()%200;
}
```

Please observe, that the numbers may not be unique.

Array Initialization

Shuffling

An array containing unique numbers arranged in an ascending order may be created by assigning to each of its elements the value of the element's index. Such an array can be changed into an unsorted array with unique numbers by applying the shuffle algorithm. The algorithm generates a permutation of the values by visiting each of the elements of the array, except for the last one, and swapping its value with other (pseudo) randomly chosen element. The latter element is selected from a set constituted of the yet not visited elements of the array and the currently visited element. This algorithm is implemented with the help of three functions that are described in the next slides.

Array Initialization

Shuffling — Swapping the Value of Elements

Below is defined a function that swaps the value of two variables. These variables are passed to the function by pointers.

```
void swap(int *first, int *second)
{
    int tmp;
    tmp = *first;
    *first = *second;
    *second = tmp;
}
```

Array Initialization

Shuffling — Choosing an Element

The function presented in this slide chooses the index of an element of the array, that belongs to the set consisting of the currently visited element and the yet not visited elements. The index of the currently visited element is passed to the function by the `from` parameter. The number of elements of the array is determined by the `length` parameter.

```
int choose(int from, const unsigned int length)
{
    return from+rand()%(length-from);
}
```

Array Initialization

Shuffling — the Implementation

The function defined below shuffles the values in the array. Please notice, that the `choose()` function is invoked in brackets. It means that the value returned by the function is used as the index of the element whose value is swapped with the value of the currently visited element of the array. The index of the latter one is in the `i` variable.

```
void shuffle(int array[], const unsigned int length)
{
    srand(time(0));
    unsigned int i;
    for(i=0;i<length-1;i++)
        swap(&array[i], &array[choose(i, length)]);
}
```


Making Copy of Array

Two arrays of the same number and the same type of elements may be copied with the use of a loop. However, the task may be accomplished more efficiently using the `memcpy()` function. It is declared in the `string.h` header file. The function takes three arguments. The first one is the name of the destination array. The second one is the name of the source array. The third one is the size of the copied (source) array. Generally, the sizes of the arrays may differ. In that case the destination array should be bigger than the source array. The value returned by the `memcpy()` is usually ignored. Another useful function, that is also declared in the `string.h` file, is `memset()`. It also takes three arguments. The first one is the name of an array. The second one is an `int` number that should be copied to all elements of the array. The last argument is a size of the array. The `memset()` function could be utilized to initiate elements of the array with the same value. It is also a common practice to ignore the value returned by `memset()`.

Printing the Values of Array's Elements

An iterative statement, like the `for` loop, may be used for indexing and successively visiting all elements of an array. Below is defined a function that uses such a loop to display all values stored in an array.

```
void print_array(int array[], const unsigned int length)
{
    unsigned int i;
    for(i=0; i<length; i++)
        printf("array[%u]: %d ",i,array[i]);
}
```

The values may also be displayed in a simpler manner:

```
void print_array(int array[], const unsigned int length)
{
    unsigned int i;
    for(i=0; i<length; i++)
        printf(" %d ",array[i]);
}
```

Thanks

Many thanks to Grzegorz Łukawski, PhD and Leszek Ciopiński, MSc for helping me to complete the Polish version of this slides.

Questions

?

THE END

Thank You for Your attention!