

UNIVERSITY OF TWENTE.

Faculty of
Engineering Technology

Programming in Engineering

Stefan Luding, Anthony Thornton

edition: 2010

Contents

Contents	i
Preface	vii
Acknowledgement	vii
History	vii
A C++	1
1 Introduction to C/C++	2
1.1 Objective of the course	2
1.2 Installation of a C++ compiler/GUI	2
1.2.1 Windows XP	3
1.2.2 Windows VISTA/ Windows 7	3
1.2.3 MAC	4
1.3 A short history of C/C++	4
2 Programming-style	5
2.1 Comments (see also §3.1)	5
2.2 Clear style	6
3 The basics	7
3.1 Language Elements	7

3.2	Data-types, Variables, Constants	10
3.2.1	Numbers	10
3.2.2	Basic data-types and their declaration	11
3.2.3	Variable-attributes	13
3.2.4	string type	13
3.2.5	String stream	14
3.2.6	Life-time and position in memory	15
3.2.7	Fields	16
3.2.8	Valarray	17
3.3	Summary	18
4	Decision structures	20
4.1	The if-command	20
4.2	The ternary ?: operator	21
4.3	switch-command	22
5	Loop control structures	23
5.1	(do) while-loops	23
5.2	for-loops	24
5.3	break and continue	24
5.4	goto-command, labels	25
6	Functions and Operators	27
6.1	Transfer of arguments	29
6.1.1	Example of passing by reference	29
6.2	Functions without arguments or return value – void	30
6.3	Libraries	31
6.4	Operators	32
6.4.1	Arithmetic operators	32
6.4.2	Comparison and logical operators	33
6.4.3	Further operators	34

6.5	inline functions	34
6.6	Overloading of functions	34
7	Pointers, pointer-field duality, and references	37
7.1	Pointers	37
7.1.1	Pointer to void	40
7.2	pointer-field duality	40
7.3	Transfer of field-variables to functions	41
8	I/O (Input and Output)	45
8.1	Elementfunctions of iostreams	45
8.2	Formatting	46
8.3	Files	47
9	Some advanced topics	49
9.1	Dynamical Memory Management	49
9.2	Object Oriented Programming	50
9.2.1	Example of Object Oriented Code involving Classes	50
9.3	Organisation implementation and header-files	52
9.3.1	Compiling and linking	53
10	Literature and longer exercises	54
10.1	Further reading	54
10.2	Exercises C/C++	55
B	MATLAB	60
11	Getting started with MATLAB	61
11.1	Input via the command-line	61
11.2	help-facilities	64
11.3	Interrupting a command or program	65

11.4 Path	65
11.5 Workspace issues	66
11.6 Saving and loading data	66
11.7 Exercises	67
12 Basic syntax and variables	68
12.1 MATLAB as a calculator	68
12.1.1 An introduction to floating-point numbers	69
12.1.2 Assignments and variables	70
12.2 Exercises	72
13 Mathematics with vectors and matrices	73
13.1 Vectors	73
13.1.1 Colon notation	74
13.1.2 Extracting and appending parts of a vector	74
13.1.3 Column vectors and transposing	75
13.1.4 Product, divisions and powers of vectors	76
13.2 Matrices	78
13.2.1 Special matrices	80
13.2.2 Building matrices and extracting parts of matrices	81
13.2.3 Operations on matrices	83
13.3 Exercises	85
14 Scripts	89
14.1 Script m-files	89
14.2 Exercises	90
15 Visualization	91
15.1 2D plots	91
15.2 Several functions in one figure	92
15.3 Adding text	94

15.4	Editing plots	94
15.5	Changing the axis	95
15.6	Exporting graph	96
15.7	Plotting surfaces	96
15.7.1	Contour plots	97
15.8	3D line plots	98
15.9	Animations	98
15.10	Exercises	99
16	Control flow	101
16.1	Logical and relational operators	101
16.1.1	The command <code>find</code>	103
16.2	Conditional code execution	105
16.2.1	Using <code>if ... elseif ... else ... end</code>	105
16.2.2	Using <code>switch</code>	107
16.3	Loops	108
16.4	Evaluation of logical and relational expressions in the control flow structures .	110
16.5	Exercises	111
17	Functions	113
17.1	Function m-file	113
17.1.1	Subfunctions	116
17.1.2	Special function variables	116
17.1.3	Local and global variables	117
17.1.4	Indirect function evaluation - <i>optional</i>	117
17.2	Scripts vs. functions	119
17.3	Exercises	119
18	Writing and debugging MATLAB programs	122
18.1	Structural programming	122
18.2	Debugging	124

18.3 Recommended programming style	125
19 Cell arrays and structures - <i>optional</i>	127
19.1 Cell arrays	127
19.2 Structures	128
20 File input/output operations	131
20.1 Text files	132
20.2 Working with Excel	134
21 Numerical analysis	135
21.1 Differentiation	135
21.2 Integration	137
21.2.1 MATLAB commands	137
21.3 Solving differential equations - The ODE toolbox	139
21.3.1 1st order ODE	140
21.3.2 Systems of ODE	141
21.4 Exercises	142
22 Exercises Day 1	144
22.1 Matrices and Vectors	144
22.2 Visualization	146
23 Exercises Day 2	148

Preface

The manuscript contains a short introduction into the interpreter language MATLAB and the object-oriented compiler language C++ extended by different small and larger exercises.

It consists of two parts : First part an introduction to C++ and second part MATLAB . Both parts are self-contained, but cross-references and similarities are highlighted.

Therefore, students who are interested in self-study, could begin either with the C++ or the MATLAB part. Within both chapters, you will find small examples. These examples can be used for self-evaluation. At the end of the C++ and the MATLAB chapter, you will find various larger exercises similar to the exercises of the final assignment.

Acknowledgement

This work includes material from “Programming in Engineering” by Stefan Luding, Holger Steeb & Katja Bertoldi (University of Twente) and “Introduction to Matlab” by Elżbieta Pękalska (Delft University of Technology) with contributions from Hans Zwart (University of Twente). The respective authors have given us permission to re-use their work.

History

Version 1.0 The manuscript was compiled for the 10-day summer course *Programming in Engineering*, July - 18 July 2008.

Version 1.1 Updated for self-study purposes, 3 September 2008

Version 1.2 Compiled for the 10-day summer course *Programming in Engineering*, 6 July-17 July 2009. This version also used for self-study program beginning 2 September 2009.

Version 1.3 Extended and updated by Stefan Luding and Anthony Thornton, for the 10-day summer course *Programming in Engineering* 5 July - 18 July 2010. This version also will be used for self-study program beginning 30 August 2010.

September 13, 2010

Stefan Luding
Anthony Thornton

Part: A

C++

Chapter 1

Introduction to C/C++

1.1 Objective of the course

The goal of this short course in C/C++ is not to become a perfect C/C++ programming guru. For this we refer to special courses and a broad range of expert literature – see below in section

The goal of this introduction to C/C++ is to better understand computer hardware and software, to perform basic tasks, operations, and programming within the framework of the programming language C/C++. Goal of this course is to first install yourself an advanced Graphical User Interface (GUI) software toolbox like DEV-C++ and make you understand simple C/C++ programs, allow you to write own code, and to get used to error-searching/debugging of code.

First of all you need a C++ compiler. On UNIX, LINUX and MAC systems this should be installed already. (If not, install the `g++` compiler and debugger.) On MS-Windows you have to follow the instructions in section 1.2.

1.2 Installation of a C++ compiler/GUI

There are many commercial C++ compilers for various operating systems available. While many Unix-type operating systems like modern Linux distributions have C++ compilers and debuggers embedded (`g++` and `gcc`), this is not the case for a MS-Windows systems like Microsoft Windows XP or Microsoft Vista. Nevertheless, you can use the following open source GUI-based C++ compiler (GNU General Public License (GPL)) on a MS-Windows system.

1.2.1 Windows XP

- **Download.**

Go to: <http://sourceforge.net/projects/dev-cpp/> and click on:
Download Dev-C++. Download the Binaries package devcpp-4.9.9.2_setup.exe.

- **Installation.**

Double click on the setup file. On the **Choose components** window, select the **Full** type of installation.

- **First time configuration.**

Accept all the default configurations.

A very useful tool is to have line numbers displayed.

Tools → Editor options → Display → Line numbers

In case the linker does not find some functions, you have to set the following command line options.

Tools → Compiler options

and insert `-lm -l .` into both fields and activate the input.

- **Ready to start.**

First create a new C++ source file and save it under the name e.g. `test.cpp`.

```
1 #include <cstdlib>
2 #include <iostream>
3 using namespace std;
4 int main(int argc, char *argv[])
5 {
6     system("PAUSE");
7     return EXIT_SUCCESS;
8 }
```

Alternatively, you can create a new project:

File → New → Project → Console Application.

At this point the project is created with a main file in it. You need to save the main file under the name e.g. `test.cpp`. Then you have to leave the project and open the source file (projects are much too complicated/overshooting for our purposes - we recommend to avoid them). Then you can compile it and run it!

Furthermore, you can try to modify the file and compile and run it.

```
1     cout <<"hello!"<< endl;
2     system("PAUSE");
3     return EXIT_SUCCESS;
```

1.2.2 Windows VISTA/ Windows 7

Here is the solution to install Dev-C++ under windows Vista.

<http://www.techsupportforum.com/pro/programming/186469-good-free-vista-compatible-c-compiler.html>

Re: Good, Free, Vista-Compatible C++ Compiler

I had the same problem, looked around, and found one that really works great for me. It's the Dev-C++, however, it's the beta version of one. The regular one didn't work for me. This is what it's called:

Dev-C++ 5.0 beta 9.2 (4.9.9.2) (9.0 MB) with Mingw/GCC 3.4.2

you can find it here, just scroll down a little, should be the first one under downloads, hope that helps

<http://www.bloodshed.net/dev/devcpp.html>

1.2.3 MAC

For the MAC you need to download and install Xcode. This comes with g++ compiler and a development environment called Xcode. Unfortunately you have to register with apple to be able to download the package (only takes ten minutes). For more information

<http://developer.apple.com/technologies/tools/xcode.html>

1.3 A short history of C/C++

C was developed in the 70's by K. Thompson, B. Kernighan, and D. Ritchie, and was introduced together with the operating system UNIX. Today, LINUX is the free derivate of UNIX, and like its ancestor is based on C, i.e., the operating system is – in big parts – written in C. The original goal was a standardized language for the structured programming. However, the so-called ANSI-standard was only introduced as late as 1988 as ANSI-C.

C++ can be seen as an extension of C, developed, propagated and introduced between 1983 und 1985 by B. Stroustrup. Besides the goal to develop a higher, object oriented language, down-ward compatibility with C was desired. Therefore, C is to be seen as a fraction/part of C++. The ANSI Standard for C++ was only fixed in June 1998. In the framework of this course, we will use C++ syntax whenever this is simpler than C, e.g. for I/O (input/output).

Some reasons for the use of C/C++ are

- it is free! (e.g. g++ compiler or DevC++ environment)
- it is system-independent – applicable also on super-computers
- it is rather low level – close to the operating system
- it is extendable, re-usable, and robust, ...

Chapter 2

Programming-style

This section briefly discusses some issues related to a good style of programming. Good style means that programs are simple and easy to read. It involves some rules like indentations that make the program also optically structured and, finally, it involves comments that describe in words what the program is doing.

Note that a good style does not only help other people reading and understanding your program, it also will help you to understand your own program – after you have not looked at it for a while. Think not of the few 100 lines we will write in this course, think of several 10,000 lines of code that a programmer can write during a thesis for example.

In “real life” people do NOT spend most of the time with programming itself, but with **debugging** (error-search), updates, changes, optimization and re-using of existing codes and code-fragments. For these tasks, a clean, proper, and consistent style is not only helpful, but absolutely necessary.

This is – by the way – true for all programming languages, including MATLAB.

2.1 Comments (see also §3.1)

There are two things a program should do: First, it tells the computer what to do and, second, it also tells the (human) reader of the program what it does.

A working program without comments is a time-bomb. Even the simplest ideas flowing into the program will be forgotten after some time ... and without comments, a program can not be re-used. Without comments, a program cannot be understood or changed so that either it has to be program anew, or a lot of time is spent re-understand the exciting program; both are a waste of time.

In C++ there are two types of comments.

The first, traditional comments begin with `/*` and end with `*/`. They can enclose single characters or many lines and are most convenient for larger comments, or for making old code invisible to the compiler.

The alternative comment begins with two slash `//` and ends automatically at the end of a line. It can start anywhere in the line and is thus convenient for commenting on single commands.

As a suggestion, each program should contain at the beginning:

- ***Header***
Name of program and what is it supposed to do.
- ***Author***
Author and (e-mail) contact address
- ***Date***
Date of changes
- ***How to use the program***
Here the mode of usage is described
- ***Files (I/O)***
Input data file format and output data-files
- ***Restrictions***
Are there, e.g., restrictions to the input-data-files?
- ***Revisions***
Dates and type of changes performed
- ***Error-management***
What happens when a program detects an error
- ***Other remarks***
...

2.2 Clear style

A program should be read like a thesis or a scientific paper. It should contain an introduction, like the header-comments above, and it should be clearly separated in chapters, paragraphs, etc. with clear boundaries. Comments can be used as a way to optically separate paragraphs within the program.

Chapter 3

The basics

3.1 Language Elements

The first language element are *comments*. They do not affect the function of the program or the computer – they enhance the understanding of the program when read by another (human) user. A comment is everything between `'/*'` and `'*/'` or after `'//'`.

The second essential element are *data*, i.e., counters, physical constants, variables, etc. It is necessary to first tell the computer which data are required for the program to work. This could be done anywhere in the program, however, it is neither efficient nor clean or clear programming. It is recommended to define data (only) at the beginning of a function/program. Exceptions to this rule will be discussed below.

Operators, *Control Structures*, and *Functions* are required to manipulate the data. Examples for operators are the basic computations `+`, `-`, `*`, or `/`. However, there are many more as summarized in the following table and discussed in more detail later in this script.

expression	short-form for:
<code>x += y</code>	<code>x = x+y</code>
<code>x -= y</code>	<code>x = x-y</code>
<code>x *= y</code>	<code>x = x*y</code>
<code>x /= y</code>	<code>x = x/y</code>
<code>x %= y</code>	<code>x = x%y</code>

Table 3.1: Short expressions and the long form they replace

Examples for control structures are `for`- and `while`-loops, situation dependent execution with `if`, `else if`, and `else` structures, block begin- and end-brackets `{ ... }` and many others as defined in the next subsections.

operator	evaluation from
::	left to right
() [] -> . X++ X-- typeid XXX_cast<TYPE>()	left to right
! ~ ++X --X + - * & (TYPE) sizeof new delete	right to left
.* ->*	left to right
* / %	left to right
+ -	left to right
>> <<	left to right
< <= > >=	left to right
== !=	left to right
&	left to right
^	left to right
	left to right
&&	left to right
	left to right
= *= /= %= += -= <<= >>= &= = ^=	right to left
?:	right to left
throw	—
,	left to right

Table 3.2: Priority and associativity of operators

Examples for functions are the trigonometric functions `sin` and `cos`, but functions can also be user-defined groups or sequences of operations that are repeated many times. Use of functions makes the program shorter and better readable.

Note that the so-called main program `main()` is also a function – it is *always* called `main`. (That means no other function can use this name.) Begin and end of a function are marked by the curly brackets ‘{’ and ‘}’.

The next class of language elements are those used for the communication between the user and the computer, the so-called I/O (input/output) functions and operators. In C/C++ normed libraries exist already that contain the definitions of the I/O commands.

The simplest way of I/O is output on the screen and input via the keyboard. The next way – especially when a lot of information has to be input into the program or is returned – is via data-files (for both input and output).

Finally, there are *organisational* elements, that do not per-se belong to the programming language, but can be used for management of the program *before* compilation and execution, e.g., by the so-called pre-processor. This allows to shorten and clarify the program by just moving big, un-clear code-fragments out of the program, or by conditionally selecting alternative parts of the code.

The first program (DevC++ New Program) looks like:

```
#include <cstdlib>
#include <iostream>
```

```
using namespace std;

int main(int argc, char *argv[])
{
    system("PAUSE");
    return EXIT_SUCCESS;
}
```

- The first lines include libraries `#include <iostream>` and `#include <cstdlib>`, get the definitions from the standard libraries for I/O and standard C-routines and functions. The former are required for input and output of text and numbers, the latter contain various helpful and necessary tools. This line must be situated at the beginning of each main-program – outside of the `main()...` function. This library contains operators/functions like `'cin'`, `'cout'`, `'<<'` or `'>>'`.
- The command `using namespace std` imports all symbols (like `cin` or `cout`) from namespace `std`. Such namespaces in C++ are used, to separate algorithms from different libraries. A command `cout` from another namespace, e.g., a self-defined one, cannot be used anymore. This is supposed to help against confusing names for functions: in principle, C++ allows “overloading”; even though this provides an enormous freedom, it is better to avoid it in order to avoid confusion. In general, we will use the C++ standard libraries, whereas for using C-functions, the command `using` is explicitly needed.

Note: You can also import individual symbols, for example `using std::cout` imports the command `cout` from the `std` namespace. We will see what this command does in §4.1

- The command `main(...)` is a function of type Integer `int`. Every program contains at least one such function (there can be much more functions, but only one special “main”). The terms in brackets are used to transfer parameters to the function `main`, however, we delay the discussion of this to later.
- `return EXIT_SUCCESS;` returns the value of 0 to the calling program, i.e., to the operating system. `EXIT_SUCCESS` is defined in the standard library as integer with value zero. Zero is the common return value if the function finished without errors.
- The brackets `{...}` enclose the main-program, i.e. the commands, the main function `main(...)` consists of. In general, program blocks are enclosed by these curly brackets, and each begin-bracket `{` must be accompanied by an end-bracket `}`.
- The command `system("PAUSE");` writes the word “Pause” to the screen and then waits for a key-stroke from the user. This is a convenient way to stop the program before it completely ends and disappears from the screen.

Note : This is a windows only command, in linux or on a MAC your program will run in a terminal anyway so the final output can always be seen.

- Note that *every* command has to be ended by a semicolon `;`. This is because a new line alone is not the beginning of a new command or program element. It is better to not use several `;` within a single line in order to keep the program clean and readable. Since the semicolon has the function to separate two commands, it is *not* found at the end of program-blocks, that are already clearly terminated by the end-bracket of the pair `{...}`.

3.2 Data-types, Variables, Constants

3.2.1 Numbers

Computers process numbers (and also symbols) in binary representation. The binary system was developed by G.W. Leibniz when studying Chinese letters. Based on the binary representation, also octal (3 bits) or hexa-decimal (8 bits) representations are used in the computer.

Example – if you really want to know:

The number 496 can be represented in the following ways (where the subscript indicates the type of representation, and the superscript are mathematical powers):

$$\begin{aligned}
 496_{10} &= 4 \cdot 10^2 + 9 \cdot 10^1 + 6 \cdot 10^0 \\
 &= 1 \cdot 2^8 + 1 \cdot 2^7 + 1 \cdot 2^6 + 1 \cdot 2^5 + 1 \cdot 2^4 + 0 \cdot 2^3 + 0 \cdot 2^2 + 0 \cdot 2^1 + 0 \cdot 2^0 \\
 &= 111110000_2 \\
 &= 7 \cdot 8^2 + 6 \cdot 8^1 + 0 \cdot 8^0 \\
 &= 760_8 = 0760 \quad (\text{in C++}) \\
 &= 1 \cdot 16^2 + 15 \cdot 16^1 + 0 \cdot 16^0 \\
 &= 1F0_{16} = 0x1F0 \quad (\text{in C++})
 \end{aligned} \tag{3.1}$$

The number 5 is much simpler:

$$\begin{aligned}
 5_{10} &= 5 \cdot 10^0 \\
 &= 1 \cdot 2^2 + 0 \cdot 2^1 + 1 \cdot 2^0 \\
 &= 000000101_2 \\
 &= 0 \cdot 8^1 + 5 \cdot 8^0 \\
 &= 5_8 = 0005 \quad (\text{in C++}) \\
 &= 0 \cdot 16^1 + 5 \cdot 16^0 \\
 &= 005_{16} = 0x005 \quad (\text{in C++})
 \end{aligned} \tag{3.2}$$

Numbers x in floating-point representation ($0.496\text{e}+03$) are transformed to fixed-point notation using the formula

$$x = mb^e \tag{3.3}$$

with the “mantissa” m (here 496), the basis b ($b = 10$ could be also $b = 2, 8, 10, 16$), and the exponent e (here 3). The in-fact implementation of floating-point numbers is machine/processor-dependent and will not be discussed here.

Exercise #2 is related to this paragraph ...

3.2.2 Basic data-types and their declaration

The most basic built-in data-types in C++ are: `bool`, `char`, `int`, `float`, `double`. Historically the processors were optimized to deal with those – however, the development of processors has progressed so much that also other data-types are efficiently treated. Usually, a compiler can optimize a code such that a user does not feel anymore the difference between different types of data.

The type `bool` is a logical data-type that can be used, e.g. to deal with the result of a logical comparison using logical operators, see Sec. (6.4). It can only have two values `true` or `false`. Furthermore, its freely transformable to data-type `int`, where `true` corresponds to 1 and `false` to 0. This transformation was introduced for compatibility reasons, to old programs where instead of `bool` just `int`'s was used. Some simple declarations are:

```
bool    test;           // declare test
test   = (0 < 5);       // comparison with result true
                        // test is set to true

int     i(test);        // i is declared and set to 1
int     j=0;            // j is declared and set to 0
```

Initialisation/Construction of a variable can be performed as *type variable = type_constant* or equivalently as *type variable(type_constant)*. (For experts: The equal symbol '=' here is not the setting of the variable, since this was just created; technically, here is *not* the set-to Operator '=' called, but a copy-operator.)

The type `int` is that type that is recognized by the compiler (i.e. the compiler-programmer) as most appropriate for dealing with integer numbers. Its size is not fixed in the standard. However, with help of the operator `sizeof`, one can determine the size of a type `int` in units of a data-type `char`, which is also the smallest data-type (again, selected by the compiler as most appropriate for single letters).

Different computers also support further integer data-types, e.g., 16, 32, or 64 bits large. These are defined using `short` and `long` or (not ANSI conform) `long long`. The longer form of declaration `short int` and `long int` is not required. Note that an integer variable always has an integer value – even if one assigns a real (floating-point) value to it.

Examples for integer declarations are:

```
bool    test;           // declare test
test   =(0 < 5);        // (see above)

int     i(test);        // i is declared and set to 1
int     j=0;            // j is declared and set to 0
long    l, n, m;        // l, n, and m are declared "long"
short   s1, s2;         // s1 and s2 are declared "short"

cout << sizeof(test) << " "
    << sizeof(i) << " "
    << sizeof(l) << " "
```

```

    << sizeof(s1) << endl;    // check size of: test, i, l, s1

i = 37;
j = 38.5674;                // ATTENTION
cout << i << " " << j << " " << i/j << endl;
                             // this leads to output: 37 38 0

```

The data-types `float` and `double` are the single- and double-precision floating point data-types that are implemented optimally in the processor. The type-specifier `long` can be used to modify `double` in order to get quad-precision (if supported). The in-fact representation of floating-point numbers is hardware dependent, e.g., in 1998 at least the double-precision basic operations were hardware implemented. Nowadays, with 64-bit processors more and faster operations with larger numbers are possible. Type `float` should only be used if memory is a problem, since for the (processor) operations, often `float` has to be modified to `double`.

C/C++ does not specify how variables are represented in memory and the use of memory is not specified per type either. However, at least the relations:

```

1 = sizeof(char) = sizeof(bool)
≤ sizeof(short) ≤ sizeof(int) = sizeof(unsigned) ≤ sizeof(long)
and
sizeof(float) ≤ sizeof(double) ≤ sizeof(long double)  are always true.

```

In general, on UNIX-machines, `int`-variables are 32, `short` 16, and `char` 8 bit long. Double-precision `double` requires 64 bit, and `float` only 32 bit. The values that are used are different from machine to machine, but also can be different from compiler to compiler.

In the following example, different data-types are declared:

```

#include <string>
using string;                // Define this at the very beginning

char    c1 = 'x';           // one character: 'x'
char    c2[] = "hello";     // text array with 5 characters
                             // (note: length is 6 due to added '\0')
string  str("hello");       // the C++ way, arbitrary length strings
                             // (note: length is 5 - no added '\0')
char    c3 = '\n';          // the new-line character
int      i_var;              // integer variable with name i_var
long int il_var = 1;         // long constant
long int jl_var = 1L;        // another long constant
short    is_var;             // short integer variable (int per default)
double   d_var = 34.2;       // real number 34.2 with double precision

```

Note that for `char` variables single characters can be enclosed by two apostrophes `'`, where as longer chains of characters must be enclosed by a double apostrophe `"`. Above, the variables `c1` and `c3` both have length unity (1), whereas the chain `"hello"` is accessible in the “field” `c2` and has length 6, since during its declaration, automatically, a symbol `'\0'` is added. (*Try the definition `char c2 = 'hello';` that will lead to an error.*) More flexible is the C++ method to

use `string`, which do not need the terminator `'\0'`. Note that for strings, the standard-library `string` must be included, see above.

For a computer there are three types of single quotes, left quotes `'`, right quotes `'` and straight quotes `'`. Characters must be enclosed by two straight quotes or your code will not compile. Some operating systems try and be clever and insert `'` the first time you ask for a quote and `'` the second time. If you get this problem turn off sticky keys on your operating system and it should disappear.

Declaration of variables is not limited to the beginning of a block, but can be done everywhere at first occurrence of a variable – be it clear, useful, and efficient or not.

Important and useful is the declaration within a `for(;;)` loop, as discussed in more detail later:

```
double a;
cin >> a;           // input a via keyboard
int    i_trunc = a;  // throw away the fractional part
for ( int i=0; i < i_trunc; i++ )
{
    ...
}
```

In the above program, the variable `i` is outside of the `for(;;)` loop not existent, i.e., neither before nor after.

Summary

- (i) Declarations can be located at arbitrary position in the program
- (ii) Initialisation with `()` or `=` form
- (iii) `sizeof(type)` can be used if the size of a type must be known, e.g., for portability between different computers/architectures.

3.2.3 Variable-attributes

Besides the attributes `short` or `long` integer variables can be also modified by `signed` and `unsigned`, that allow to use integers (positive or negative) or exclusively positive integers, respectively. The `int` can be dropped here.

Exercise 3.1.

Write a program to find out the size of variables of different types on your computer. ■

3.2.4 string type

String objects are a special type of container, specifically designed to operate with sequences of characters. Each character of the string can be accessed using the `[]` operator i.e.

```
string a;
```

```
a="Hello";  
string b;  
b=a[2];  
cout << b <<endl;
```

would write out the letter 'l' , recall c++ arrays start numbering at 0.

Additional with the + operator can be used to concatenate two strings together, for examples

```
string a,b,c;  
a="Hello ";  
b="World";  
c=a+b;  
cout << c << endl;
```

would produce 'HelloWorld'.

Note: The function `c_str()` returns a const pointer to a regular C string, identical to the current string. The returned string is null-terminated. This is often required because a lot of routines require C strings not C++ strings to be passed to them.

3.2.5 String stream

Stringstream provides an interface to manipulate strings as if they were input/output streams. It is located in the header called `sstream` as you would expect from the other stream names, but the type is called `stringstream`.

This is described here, after string for ease of finding, but `iostreams` are described in more detail in chapter 8, so please refer to this section for a more detailed explanation of streams.

Below is a small example code that creates a filename by passing information into a `stringstream` and re-extract the string at the end.

```
#include<sstream>  
#include<fstream>  
#include<iostream>  
  
stringstream file_name;  
stringstream problem_name("my_problem");  
ofstream script_file;  
file_name << problem_name.str() << ".disp";  
script_file.open((file_name.str()).c_str());
```

The key commands are `<<` which adds data to the `stringstream` this can be any type for which the `<<` operator is defined. The `.str()` method return a traditional C++ string, see section 3.2.4 were `.c_str()` is also described.

The nice thing about stringstream is they automatically deal with type conversion. For example the code below will convert a string to int via a stringstream.

```
#include <iostream>
#include <sstream>
using namespace std;
int main()
{
    int a;
    string s = "456";
    istringstream sin(s);
    sin >> a;
    cout << a << endl;
    return 0;
}
```

3.2.6 Life-time and position in memory

auto If there is no further declaration, variables are **automatic**. This means they are constructed as soon as the program reaches the line where they are declared. As soon as the program leaves the block where they are declared, the memory is freed again. (The block is, in general, the inner-most enclosing pair of brackets, or like in the case of a `for(;;)` loop, the end of this object.) *For experts: Variables of this type are usually assigned by the compiler to the so-called stack.* The value of such a variable is random! if not explicitly assigned or initialized. Thus *make sure to initialize your variables!*

Aside: stacks, heaps, etc.

In computer science, a **stack** is a last in, first out (LIFO) abstract data type and data structure. A stack can have any abstract data type as an element, but is characterized by only two fundamental operations, the push and the pop. The push operation adds to the top of the list, hiding any items already on the stack, or initializing the stack if it is empty. The pop operation removes an item from the top of the list, and returns this value to the caller. A pop either reveals previously concealed items, or results in an empty list. (See Wikipedia for further reading).

A **queue** is a particular kind of collection in which the entities in the collection are kept in order and the principal (or only) operations on the collection are the addition of entities to the rear terminal position and removal of entities from the front terminal position. This makes the queue a First-In-First-Out (FIFO) data structure. (See Wikipedia for further reading).

In computer science, a **heap** is a specialized tree-based data structure that satisfies the heap property: if B is a child node of A , then $\text{key}(A) \leq \text{key}(B)$. This implies that an element with the greatest key is always in the root node, and so such a heap is sometimes called a max-heap. (Alternatively, if the comparison is reversed, the smallest element is always in the root node, which results in a min-heap.) The several variants of heaps are the prototypical most efficient implementations of the abstract data type priority queues. Priority queues are useful in many applications. In particular, heaps are crucial in several efficient graph algorithms. (See Wikipedia for further reading).

static The keyword **static** makes a variable exist during the whole duration of the program. *For experts: Such variables are saved on the heap-memory.* If not initialized otherwise, they get the value 0 (zero). A **static** variable inside a function has thus the value 0, if the function is called for the first time, and keep this (or another assigned value) between subsequent calls to this function. In contrast, an **automatic** variable is every time newly created and initialized. The following function counts how often it was called and returns this value.

```
int f_count_calls(){
    static int count; // 0 at first function call
    return ++count;   // increment before use as return value
                    // (i.e. returns one on the first call)
}
```

const Variables that are declared **const**, can only be assigned a value at initialization and not anymore thereafter. This is used to define constants within a namespace that – by the compiler – can be optimized such that their use costs less time than using variables. Thus, if variables can be defined as **constants**, this might lead to a faster program execution.

```
const int    Size = 100;
const double Pi  = 3.1415926; // Self-defined value of PI
                                // When including cmath one can
                                // use the predefined constant M_PI
```

3.2.7 Fields

Fields are used to combine variables of the same type in a connected range of memory. *For experts: Access to variables can be faster if they are in such a connected range of memory and, on the other hand, using fields also can make a program shorter and more clear.* Access to the elements of a field is achieved by an integer index. A linear field (with one index) can be seen as a vector, while a field with two indices represents a matrix. Mathematical formulas with indices can often be represented in a program using fields.

A field is declared in C/C++ by adding brackets `[]` to the variable name. At the same time, this operator is also used to access the elements of the field, see below. In C/C++ fields always start with the index 0, so that the last valid index is the length of the field reduced by 1. Multi-dimensional fields are declared by multiple (not enfolded) pairs of brackets, see below.

Such fields receive – at declaration – a fixed size. Thus, the argument in brackets `[]`, must be either a number, a **const** variable, or a well-defined variable with a known value for the compiler (at compilation time).

```
#include <string>
using string;

// array of 20 int's
int a[20];
```

```

// initialization of field a -- ERROR!
for ( int i=0; i <= 20; i++ )
    a[i] = i;
// initialization of field a
for ( int i=0; i < 20; i++ )
    a[i] = i;

// array of three strings, initialized to some values:
string  stra2[] = { "Katia", "Holger", "Stefan" };

// two dimensional array of 20x30 elements of int type
int  ia[20][30];
// initialization of field ia
for ( int i=0; i < 20; i++ )
    for ( int j=0; j < 30; j++ )
        ia[i][j] = i*j;

```

Later, we will see that fields and pointers are closely related. The use of pointers allows to define “dynamic” fields of a size that is determined during run-time, e.g., following a user-decision and -input.

3.2.8 Valarray

C++ supports two new containers that allow simpler manipulation of fields. This first is valarray, which will be introduced by the following example.

```

#include <iostream>
#include <valarray>

using namespace std;

int main()
{
    //Declay an integer array of length ten
    valarray<int> my_vec(10);

    //Assign the value i squared to each element
    for (int i=0;i<10;i++)
        my_vec[i]=i*i;

    //Write out each element's value to the screen
    for (int i=0;i<10;i++)
        cout << my_vec[i] << endl;

    //Write out the sum of all the elements to the screen.
    cout << my_vec.sum() << endl;
}

```

All math functions declared in `cmath` can be used with `valarray` for examples `sin` or `cos`. A major advantage of `valarray` is it supports dynamic size control, i.e. by calling `my_vec.resize(20)` you can change after creation the size of `my_vec`, this is useful when the size of a field is not known at compile-time.

In the above examples `my_vec` is a class which has both data (in this case integers) and functions (e.g. `resize` or `sum`). For an example of how to create your class see §9.2.1.

`valarray` can only store numerical data, but a very similar container class `vector` can store any type, including chars and strings. This is implemented in the header `vector`.

Access to data stored, in both `valarray` and `vector` classes can be slower than pointer based method described in §3.2.7, but the use of iterators (not covered by this course) can massively improve access time for these classes.

3.3 Summary

So far we have introduced the basic element of a program: variables, and discussed the many different types that are available. Also a few other commands have been used with out explanation, (but you can probably work out what they do) like `cout`, `cin` and `for`.

In the rest of this part, we will cover how to make decisions in a program, repeat sections of program, interact with the users and how to store data in a permanent way.

Exercise 3.2.

Explain what each line of the following code does i.e. add the comments

```
#include <iostream>

using namespace std;

int main(int argc, char *argv[])
{
    int n, m;
    cout << "Enter two integers: \n";
    cin >> n >> m;

    int sum=m+n;

    cout << "Sum of the two integers is " << sum << endl;
```

```
    system("PAUSE");  
    return EXIT_SUCCESS;  
}
```

■

Exercise 3.3.

Modify the program to ask for three numbers ■

Exercise 3.4.

Modify your program to ask for both a person first name and surname. Get the program to output the persons initials, along with the sum of the three numbers they give. ■

Chapter 4

Decision structures

For implementing algorithms into a certain programming language, this language has to provide certain control-structures that allow repeated or conditional execution of parts of the program.

In C (and C++) the most-used control structures are the `if`-statement for the conditional execution, as well as the `switch()...case:-`statement, that jumps to one of many blocks of commands, based on the value of a certain variable. Finally, there is the command `goto`, that sometimes can be used to escape from (deeply nested) loops, see §5. (Actually I would not recommend ever using a `goto`, deeply nested loops can be exited by making the outer loop a function and using the `return` command, see §6).

4.1 The `if`-command

The `if`-command is used to execute parts of the program only under certain conditions. This command can be (but does not have to be) complemented by an `else`-statement, that is visited/execute in case the condition in the `if`-command is `false`. The `else`-part can consist of arbitrarily many `else if` and one `else`. As well `if`, `else if`, and `else`-branches can be single lines (terminated by a semicolon) or blocks, enclosed in curly brackets `{}`. An `else`-block always relates to the previous, actual `if`-command, i.e., also if it is part of an `else if`. In the following example, the second `if`-command is a statement by its own, outside the `else`-part of the first `if`.

```
#include<iostream>
#include<cmath>
// ...

// solution of quadratic equations  $x^2 + p x + q = 0$ 
```

```

double p=2., q=3.;           // or other values
double sol1, sol2;
double root_arg = 0.25*p*p - q; // calculate the argument for the root

if ( root_arg > 0. )
{
    // begin of if-block
    sol1 = 0.5*p + sqrt(root_arg);
    sol2 = 0.5*p - sqrt(root_arg);
}
// end if-block
else if ( root_arg < 0. )     // 'error' condition, 'else' part
    cout << "no real solution exists" << endl;
else                         // refers to last if(root_arg < 0.)
    sol1 = sol2 = 0.5 * p;    // assignment from right to left

cout << "Solutions 1,2: " << sol1 << ' ' << sol2 << endl;

if ( sol1 > 0. )
    cout << "graph cuts positive x-axis";
// ok, else branch missing

```

4.2 The ternary ?: operator

For completeness, also the ternary ?: operator is treated here. Ternary means that this operator requires three arguments and is used to conditionally evaluate parts of expressions.

```

inline int abs(int number) {
    // compute absolute value
    return number > 0 ? number : -number ;
//Equivalent form with brackets
//return (number > 0) ? (number) : (-number);
}

```

Is the term before the question-mark true, then the term before the colon is examined. Otherwise, the term after the colon will be used. The two alternatives must be of the same type. Note that here, the use of brackets is recommended – partly because the ternary operator is of low priority and also because the terms that belong together should be made clear.

```

int i1 = 3 > 4 ? 0 : 1;    // i1 is 1
int i2 = 3 * ( 3 > 4 ? 0 : 1 ); // i2 is 3
int i3 = 3 * 3 > 4 ? 0 : 1; // i3 is 0, since * binds
                          // stronger than >

```

This operator always can be avoided by using an if-command, however, it could be more clear or at least shorter in some situations. In any case it is a possibility to write un-readable code and therefore should be used sparsely.

(As final remark, the ternary operator is also embedded in the graphics program gnuplot.)

4.3 switch-command

A command related to the `if` and the `goto` commands is the `switch`-command. Dependent on the value of an integer variable, several `case`-blocks are visited. This construct is often used when there is a larger number of possible alternatives to be chosen from.

In principle, a `switch` can always be replaced by `if...else if`-constructs, but often the `switch` is more clear and helps understanding and structuring the program.

```
char c;
cin >> c;
switch( c )
{
    // begin case block
    case 'a': cout << "hello world!";
              break;
    case 'b': cout << "HELLO ";    // fallthrough intended, comment missing break
    case 'c': cout << "WORLD!";
              break;
    default:  cout << "unknown input";
              break;              // not necessary, but good style
}                                // end case block
```

The above program-segment returns in small letters "hello world!" if the letter 'a' is typed. In case c the word WORLD! is printed, but in case b the full sequence HELLO WORLD!. This strange behavior comes from the missing `break` in the block b. This behavior, namely that a execution is performed until a `break`, is called ("fallthrough"). Since this is often not desired, using this option required a comment whenever a `break` is missing. The `default`-label is optional and is visited if none of the other labels is valid in the `case`-list.

Exercises

Exercise 4.1.

Modify the program from Exercise 3.4, so that it asks the users if he/she wants to enter 2 or 3 numbers. Read in the correct number of inputs and compute the sum. ■

Exercise 4.2.

Write a program that reads in a date in the format dd/mm/yyyy format and prints out 12th of March 2010 etc... ■

Chapter 5

Loop control structures

Often you want to repeat sections of a program, the most common way is using `for`-loops. Alternatively, the rejecting `while()`- and the accepting `do{...}while()`-loops exist.

5.1 (do) while-loops

Loops that are constructed using `while` are executed until the condition in the condition part of the loop is not true anymore. Is the condition already false at first occurrence of the `while`-loop, the loop will not be executed at all.

```
char c=0;                // 'loop initialization'
while( c != 'y' && c != 'n' ) // first time c==0, so we enter the loop
cin >> c;                // can also be a block { }
```

The symbol/character variable has to be initialised such that the loop is entered. In such cases, the `do ... while(...)` loop might be more intuitive.

```
char c;
do {
    cin >> c;
} while( c != 'y' && c != 'n' );
```


5.2 for-loops

When working with fields, one often needs loops that access all elements of a field, one by one. One can program this with **while**:

```
const int Size = 20;
// ..
int a[Size];
// ..
int i=0;           // loop initialization
while ( i < Size ) // loop condition, rejecting if false
{
    a[i] = i;
    i++;           // loop control, performed after each pass
}
```

A shorter alternative used the **for(;;)** loop. The following is almost equivalent to the above, but shorter by avoiding the line with **i**, and more clear, since the begin, end, and stepsize of the loop are defined at its start.

```
int a[20];
// ..
for( int i=0; i < 20; i++ )
    a[i] = i;           // work on elements 0 .. 19
```

There are two advantages: First, the loop-beginning, -ending, and stepping is nicely combined in one line and, second, the variable **i** is only defined within the loop and not used (and not needed) outside.

5.3 break and continue

In order to exit a single, not-nested **for** or **while** loop, one can use the **break** command. The loop is terminated and the program continues directly after the loop. Thus **break** can be used to terminate loops under certain conditions, earlier than the loop implies, or it can be used to exit “endless” loops.

```
// copy input to output
while( true )
{
    char c;
    cin >> c;
    if ( cin.fail() ) break;    // end of input or something wrong
    cout << c;
}
```

The command `continue` terminates the processing of the present block and continues, dependent on the type of the loop, at the corresponding control-command (`for(;;)`), or (`(do{}) while()`).

```
#include <cmath>                // for sqrt() function

double a[20];
// ...
double sum_sqrt=0.;
for( int i=0; i < 20; i++ )
{
    if ( a[i] < 0. ) continue;    // work on next element
    sum_sqrt += sqrt( a[i] );
}
```

Note again that `continue` and `break` only exit the innermost loop-block and cannot be used to exit nested loops.

5.4 goto-command, labels

In order to leave deeply nested loops, the `goto` command can be used.

```
void f()
{
    int a[20][30];

    for( int i=0; i<20; i++ )
        for( int j=0; j <30; j++ )
            if ( a[i][j] == 42 )        // found the answer
                goto label_end_of_loop; // multi level break

label_end_of_loop:                    // label, we jump here from inside the loop
}
```

The label with name `label_end_of_loop` marks the jump-goal of the `goto`. Within a function the name can be freely chosen – terminated by a colon `:`. The problem with `goto` is, that it is not visible *from where* the jump to this label comes. Furthermore, using `goto` often requires further use – and thus further unclear structure of the program.

It is allowed to leave local blocks using `goto`, but one cannot enter a block inside an `if`-block. Neither is it possible to jump out of a function.

Since `goto`'s always can be avoided, it is better to do so. For example the `goto` in the above function could be replaced with a `return`, see §6.

Exercises

Exercise 5.1.

Write a code which reads in N numbers and computes their sum. ■

Exercise 5.2.

Write a program to compute the squares of all integers between two numbers. ■

Exercise 5.3.

Write a program that computes the factorial of an inputed number N . Recall the factorial of $N = N * (N - 1) * (N - 2) * \dots * 1$. ■

Chapter 6

Functions and Operators

Longer programs usually contain parts or blocks that are repeated a certain (large) number of times. These parts can then be defined elsewhere (as a function) and within the program are replaced by a short-cut or place-holder, i.e., the name of the function. It is wise to give the **function** a name that already implies its function.

A function (also called procedure or subroutine in other languages), has input and output variables/parameters. In the simplest case the transfer of input-parameters is done as argument, and the output as a return-value.

C++ enforces that the types of input and output are specified. In the following example, the function **power** computes from a double **base** value and an integer number **exponent** a new number in double precision that is returned to the calling function/program. The name already implies the algorithm used here.

```
// raise 'base' to the power 'exponent', with exponent being integer,
// base being double; accept also negative exponents and check ...
double power(double base, int exponent){
    double result = 1.;    // will multiply this by base 'exponent' times
    bool    neg_exp = false;
    if (exponent < 0 )      // for negative exponents we use a single
    {                       // division at the end of the function
        neg_exp = true;
        exponent = -exponent;
    }
    for (int i=0; i < exponent; i++ )
        result *= base;    // successive multiplication

    if ( neg_exp )          // have to take the inverse
        result = 1./result;
```

```
    return result;           // return the result to the calling program
}
```

Exercise 6.1.

Search for the explanation of **Flow-Chart** in, e.g., Wikipedia, and plot the Flow-Chart of the Power-Function. ■

Note that C++ already has a function `pow` that is described by
`double pow(double base, double exponent),`
 only different in the type of the exponent.

The names of the arguments are free to be chosen. Inside the function block these names are to be used. Note that like `for` loops, there is no semicolon after the function.

The keyword `return` terminates the function at the specified position, moves the value of the variable specified after it on the stack, and returns to the calling program. In the case above, the expression consists only of a single variable `result`. Round brackets are not required here. A subroutine can contain several return keywords at different places in the function block. This can be a useful tool to deal with problems/errors. Hence the `return` command can be used to exit deeply nested `for` loop to forcing the exit of the whole function.

Finally, one has to define or declare a function before it can be used in the program. This declaration is essentially the first line of the function (with semicolon). This so-called prototype tells the compiler the name of the function and the types of it and its arguments. Actually, if you careful about the order you do not need to prototype functions, but you must place the full function definition before the first call (use) of the function. This is often tricky to keep track of and prototyping is recommended.

```
double power(double base, int exponent);
```

At first definition, the command block is replaced by the single semicolon. Note that the type is required, but the name of the arguments in the function is not needed here. However, a clear program specifies variables that tell something about their purpose.

Declarations can (in principle) be repeated arbitrarily often in a program – they are a promise to the compiler that the function will be defined later. Unlike variables, functions are always global – it is not allowed to define them inside a block with limited existence. Actually function and variables can be wrapped together in a very neat package, called a class, see §9.2.1 for more details, solving this global scope problem.

A program that reads a number from the keyboard and then uses the function `power` to finally write the result could look like this.

```
#include <iostream>
using namespace std;
```

```
double power(double base, int exponent); // declaration of 'power'

int main()                                // definition of main
{
    double in;
    cin >> in;                            // read value from keyboard

    cout << "-3rd to 3rd power of " << in << ": ";
    for( int i= -3; i <= 3; i++ ){
        cout << power(in, i) << " ";      // call power and use return value
    }
    cout << "\n";

    return 0;                             // return to operating system
}

double power(double base, int exponent)    // definition of power
{
    double result = 1.;
    // ... program text as above

    return result;
}
```

6.1 Transfer of arguments

The transfer of arguments to a function takes place “by value”, i.e., before the function is called, the arguments are copied and the function uses these copies. This way, the variables that contain the values that are transferred to the function can not be changed. If this is desired, e.g., because one return value is not enough and many variables are to be changed in a function, then one either use pointers, as discussed later in section 7 or pass by reference.

If you add `&` after a variable name in the declaration the actually variable, not a copy, is passed to the function, therefore changes to that variable within the function effect the original variable.

6.1.1 Example of passing by reference

```
#include <iostream>
#include <valarray>

using namespace std;

int add(int a, int b);
```

```

int add2(int &a, int &b);

int main()
{
    int a=1;
    int b=2;
    cout << a <<" "<<b << endl;
    cout << add(a,b) << endl;
    cout << a << " " <<b <<endl;

    cout << a << " " << b <<endl;
    cout << add2(a,b)<<endl;
    cout << a << " " << b << endl;

}

int add(int a, int b)
{
    a=a+b;
    return a;
}

int add2(int &a, int &b)
{
    a=a+b;
    return a;
}

```

Exercise 6.2.

Run this program and from the output explain the difference between the functions `add` and `add2`. ■

6.2 Functions without arguments or return value – void

The keyword `void` can be used like a type, but is not really a type. It is useful to mark functions that don't have a return value. Examples are functions that just write something to the screen or to a file and functions that modify the input-variables directly.

```

void error(string message){
    cout << "error occured: " << message << "\n";
}

```

Also functions without arguments sometimes can be useful. In C++ this is implied by an empty list of arguments like for the pseudo-randomnumber generator as declared in the C library `stdlib.h`:

```
void srand(unsigned int seed);
int rand(void);
```

It is used by first selecting a sequence of random numbers by calling (once only) the initialization function `srand(.)` with a positive number. Then, for each call of `rand()` a new, integer number is returned. In order to transform the integers to real numbers, one can e.g. divide by the maximum.

Exercise 6.3.

Use the random number generator of C/C++ and later the random numbers in MATLAB to generate a vector of 10 random numbers $r_{i=1,\dots,10} \in [0 : 1]$. ■

6.3 Libraries

As seen above, declarations/prototypes are sufficient for the compiler to translate a program. Wherever called in the program, the function-call is prepared during compilation. The function itself can be then defined and “linked” to the program later. Function-definitions are found in so-called object-files (*.o) or libraries (see section §9.3 for more information of code organisation into small objects, which are compiled and linked independently).

The classical example for a library is the `cmath` file that was included at the beginning and involves functions like the `sqrt()`. The declarations are found in `cmath` itself, and the translated/compiled definitions are found in `libm.a`. A program-block that uses the `sqrt()` function can look like this.

```
double d;
cin >> d;
if ( d < 0. )
    cout << "cannot compute the square root of negative numbers\n";
else
{
    // compute sqrt(d) and print value
    cout << "sqrt(" << d << ") = " << sqrt(d) << "\n";
}
}
```

The compiler must contain the linker-option “-lm” and must know the correct path where it can find the library. As a remark, under LINUX, the compilation would look like this:

```
g++ myprog.cc -o myprog -lm
```

In DevC++ the corresponding command can be found under “Compilation”.

The short -l means “library” and m is the abbreviation for the file `libm.a`, that one gets by adding `lib`, `m` and the ending `.a`.

6.4 Operators

Operators allow us to make a program readable. Sometimes, unfortunately, they also make a program un-readable, cryptic and unclear. A term like `3 * z + 5` is easily understandable, whereas the term `plus(mult(3,z),5)` that is expressed using functions is much less clear.

There exist unary, binary, and ternary operators, see sec. 4.2. Binary operators like the multiplication- and addition-operators, have two arguments, while unary operators, like the adress-operator `&` require only one.

6.4.1 Arithmetic operators

The arithmetic operators are `+`, `-`, `*`, `/` and `%`. The Operator `%` (modulo) forms the integer rest of the division of the left with the right operand (both integer). If both are positive, the result is positive and strictly smaller than the right operand. For one or two negative operands, the result depends on implementation.

Combining the above operators with the `'='` allows to repeat the left operand, as for example:

```
int i=3, j=7;      // initialize
i    = 3 + j % i;  // i = 3 + (j mod i),   i is set to 4
i   += 3 * 4 + j;  // i = i + (3 * 4 + j), i is set to 23
i   /= 5;          // integer division, i is now 4
i   /= 5;          // integer division, i is now 0 !
```

The type that is returned by an operator expression depends on the type of the operand. If both operands are type `T`, also the result is type `T`. For different operand-types the “better” type wins; one operand is first transformed to the “better” type, then the operation is carried out. C/C++ has the following “type promotions”:

```
bool -> char -> short -> int -> long
      -> float -> double -> long double
signed -> unsigned (!)
```

These rule of type-transformation are also valid for the following operators.

The unary operators `++` and `--` exist in postfix and prefix-form, and they lead to an increment or decrement by one, respectively. In the postfix form, the value of the operand *before* the operation is returned, whereas in the pre-fix form, the value *after* is returned.

```
int i=5, j;
j = i++;          // j is 5 now (i before increment) i is incremented to 6

double a[20];
i=j;              // i=j=5
```

```

cout << a[++j] << a[j++];    // prints a[6] twice,    j is now 7
cout << a[i++] << a[++i];    // prints a[5] and a[7], i is now 7

// j++ = i;        // ERROR: cannot assign to a temporary
// i = (j++)++;    // ERROR: cannot apply ++ to a temporary

```

6.4.2 Comparison and logical operators

For operations between two integral datatypes, there exist bitwise, logical, and shift operators. (C++ also provides keywords for these). These operators are summarized in table 6.1, and they also can be combined with the =.

	bitor	bitwise OR
&	bitand	bitwise AND
^	xor	bitwise XOR
<<		left-shift, <i>n</i> -times
>>		right-shift, <i>n</i> -times
~	compl	complement, bitwise NOT, unary
=	or_eq	bitwise OR, to-left assignment
&=	and_eq	bitwise AND, with assignment
^=	xor_eq	bitwise XOR, with assignment
<<=		left-shift, <i>n</i> -times with assignment
>>=		right-shift, <i>n</i> -times with assignment

Table 6.1: Bit-manipulation-operators

The logical operators are `&&` (or **and**) and `||` (or **or**) and the negation `!`. The operators `&&` and `||` have the special property that the evaluation of this expression is terminated already if the result is clear (so-called sequencing). This property is also shared by the comma-Operator `,` – that sometimes is used to construct complicated **for**-loops. Its value is the same as the value of the right expression.

```

double    a[20];
unsigned ind[5];

// safe, even if some ind[i] >= 20, since the last expression
// will anyway not be evaluated in that case
for (int i=0; i < 5 && ind[i] < 20 && a[ind[i]] >= 0; i++ )
    sqrt(a[ind[5]]);

// sequence operator used to combine two expressions
int i, j;
for ( i=0, j=2; i < 18; i++, j++ )
    a[i] = a[j];

```

The (arithmetic) comparison-operators are `==`, `!=`, `<`, `<=`, `>` and `>=`. The value of a logical operation and the result of a comparison in C++ are of type `bool`, i.e. either **true** or **false**.

6.4.3 Further operators

The only ternary operator, `?:`, which is used to select alternative expressions of the same type, was already discussed (see sec. 4.2).

```
int i;
cout << (i > 0) ? i : 0;    // never prints a negative number
```

The operator `()` leads to a function-call, `[]` is the index-operator for fields and pointers, a point `'.'` allows to select an element from a structure and `->` is equivalent when a pointer to a structure is given.

```
class Person {
public:
    string name;
    int age;
};

void f(){
    Person    p;
    Person *p_p = &p;
    p.age     = 25;
    p_p->age = 25; // equivalent
}
```

The unary operators `*` and `&` are used to find the value of the memory location where a pointer points at, and to find an address of an arbitrary variable (see sec. 7).

Operator-expressions are (like in C) treated according to their priority and associativity, see table 6.2.

6.5 inline functions

In general, several things happen when a function is either called or closed. Some of this can be avoided or optimized by use of so-called **inline**-functions. However, we will not go into detail here. (Further reading if needed).

6.6 Overloading of functions

In contrast to C and some other languages, in C++ not only the name and type of a function, but also the types of the arguments are relevant. Therefore it is possible to define a function for each type `int`, `double` or `struct Date {...}`, all with the same name. The name does not

operator	evaluation from
::	left to right
() [] -> . X++ X-- typeid XXX_cast<TYPE>()	left to right
! ~ ++X --X + - * & (TYPE) sizeof new delete	right to left
.* ->*	left to right
* / %	left to right
+ -	left to right
>> <<	left to right
< <= > >=	left to right
== !=	left to right
&	left to right
^	left to right
	left to right
&&	left to right
	left to right
= *= /= %= += -= <<= >>= &= = ^=	right to left
?:	right to left
throw	—
,	left to right

Table 6.2: Priority and associativity of operators

have to contain type information thus. In C++, the compiler selects the right function with the appropriate type by itself. One can think of an internal function-name used by the compiler that also contains the argument-types.

```

void swap(int &a, int &b)    // internal name e.g. swap_int_int
{
    int tmp = a;
    a = b; b = tmp;
}

void swap(double &a, double &b)    // swap_double_double
{
    double tmp = a;
    a = b; b = tmp;
}

void f()
{
    double a, b;
    int    i, j;
    // ...
    swap(a, b); // double version
    swap(i, j); // int    version
}

```

Exercises

Exercise 6.4.

Write a program that reads in ten numbers, using the swap function above sort the numbers into ascending order and write out the order list. ■

Exercise 6.5.

Write a function, which calls itself, to calculate the N^{th} Fibonacci number. Recall $fib(N) = fib(N - 1) + fib(N - 2)$ and $fib(1) = 0$, $fib(2) = 1$. ■

Chapter 7

Pointers, pointer-field duality, and references

7.1 Pointers

From the viewpoint of the computer, every memory-cell in the (virtual) memory is uniquely marked by a certain number (address). C/C++ allows to use these **addresses** in memory-space for programming purposes. Since it is allowed to self-define names for these address-pointers, the program can also become much more clear and readable. The declarations:

```
int    i=5;
double a=12.34;
char   *c1="hello";    // declaration of an array of char
int    j=0;
```

create the memory-entries in Fig. 7.1 (left column). Only the memory-entries (the values of the variables) occupy space in the computer memory. The names of variables do not occupy space even if one would use instead of `i` the more meaningful name `counter`.

In general, there is no guarantee that sequentially defined variables also occupy memory sequentially (except for fields). The variable `c1` of type `char` contains a pointer to another range of memory where the values of the field are stored. The definition of `*c1` reserves some range in memory that is large enough to store the word "hello". The elements 'h', 'e', 'l', 'l', and 'o' can also be accessed using the form `c1[0]`, `c1[1]`, `c1[2]`, `c1[3]`, and `c1[4]`. Note again that numbering in C/C++ always starts with 0.

If one has to insert or delete a range of a sorted list (quickly and efficiently), then this is not possible using a linear field like the above `*c1`. A *linear field* is special in so far that all data are stored in subsequent positions in memory. As well inserting as removing entries from this field requires copying the full range "above" the modified element. (In average that is half the field-length to be moved). This is acceptable for short fields but not for long ones. A possible

value	address	(comp.def.) name	(user-def.)
5	0xfd10	i	
12.34	0xfd18	a	
0xff10	0xfd20	&c1	
0	0xfd28	j	
	...		
h	0xff10	c1[0]	
e	0xff11	c1[1]	
l	0xff12	c1[2]	
l	0xff13	c1[3]	
o	0xff14	c1[4]	
\0	0xff15	c1[5]	
	...		

Figure 7.1: (Schematic) memory occupations (left), computer-internal memory address (mid), and (user-defined) name of variables (right).

solution to this is a *linked list* which contains not only the values of the elements but also the information where the next element can be found. The actual location does not matter, i.e., only the names of the memory-cells have to be modified. For example, in order to remove an element x from a linked list, only the name entry of the previous element has to be changed such that it does not point to x anymore, but to the following element. This information is available in the x -element and thus just has to be copied. Deleting several subsequent entries requires the same effort and entering a (part of a) linked list from somewhere else is also rather efficient. However, due to this additional options, a linked list also usually requires more memory than a normal field.

In order to allow a computer independent dealing with such objects like linked lists, C/C++ provides also the type pointer. A pointer can point to an arbitrary data-type – which actually could be a pointer itself. In order to declare a pointer, a (star) $*$ is used. In order to get the adress of a variable, the $\&$ -operator is used. The declarations in the program-fragment:

```
// Declaration part
int    i = 5;           // initialisation
int  *p_i = &i;         // p_i is a pointer to an int,
                        // here the variable i
int **pp_i = (&p_i);   // pointer to pointer to int

// How to use * and & legally in the program
    i = 0;              // set i to 0
    *p_i = 1;           // set i to 1, p_i remains unchanged
```

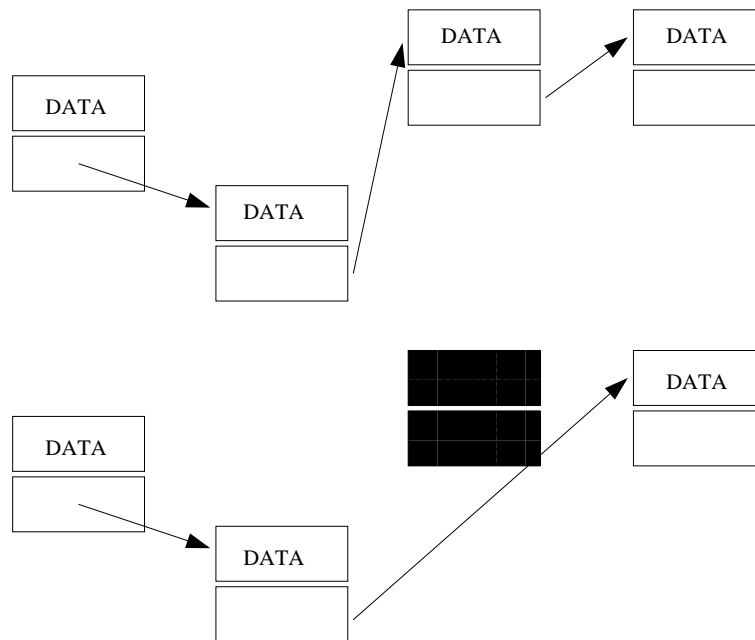


Figure 7.2: Use of pointers in a linked list. (Top) the list before an element is removed, (bottom) the list after, where the black element denotes the removed one that is not part of the list anymore..

```

**pp_i = 2;           // set i to 2, p_i remains unchanged
int j = 7;            // declare another integer j
**pp_i = &j;          // change p_i to point on j, *p_i != i
p_i = &i;             // let p_i point to i again

```

lead to memory occupation as indicated in Fig. 7.3.

5	0xfa10	i
*	...	
0xfa10	0xfc28	pi
*	...	
0xfc28	0xfe08	ppi
	...	

Figure 7.3: Memory occupation as in Fig. 7.1.

In order to change the value of a variable a pointer points to, or in order to use this value in an operation, also the `*` operator is used.

```

*p_i = 27;           // i is 27 now
*p_i = 2 * (*p_i) + 5; // equivalent to: i = 2 * i + 5;

```


Constant pointers are declared by using a `const` right of the pointer symbol `*`. They must be initialized already at declaration, since they are fixed thereafter and cannot be changed anymore.

```
const int *cp_i = &i;           // ok, i cannot be changed
                                //      through use of cp_i
                                //      cp_i can be changed

const char newline = '\n';
const char * const p_nl = &newline; // const pointer to const char
                                // note: p_nl = '0'; -> error!
```

The `*` operator has thus two meanings: In a variable declaration, it marks a pointer, whereas in front of a variable (of type pointer) it leads to access on the memory (value) of the corresponding memory space.

Analogously, also the `&` operator has two meanings: During variable declaration, it denotes a reference, whereas in front of a variable, it is the so-called **address** of operator, i.e. it provides the address of the variable.

Note that it is possible to do arithmetics/mathematics with pointers. Above, `p_i = p_i + 5` would let the pointer point to 5 `int` memory entries further. For this, the compiler needs to know how big the corresponding data-type is. Also the difference of pointers can be formed but that makes only sense inside a field.

Navigation within a field can be done in the following way:

```
int p[30];           // declare a field of int
*(p_i + 25) = 15;    // the value of the 25th entry (past the one
                    // pointed to by p_i) is set to 15
p_i[25] = 15;        // equivalent ...
```

The language guarantees that 0 is never the address of a real data-element. A pointer with value 0 can thus be used to indicate an error or, e.g., the end of a list.

7.1.1 Pointer to void

The type pointer to void is generic. It is guaranteed that any pointer can be transferred into a pointer to `void`, without loss of information. In C the `void *` was often used to transfer information to functions/routines that can deal with arbitrary pointers. (e.g. sort-function like `quicksort`, or the return value of `malloc`). In C++ it is recommended to use `template-`functions that guarantee type-consistency and type-safety. (Further reading if needed).

7.2 pointer-field duality

Between pointers and fields exists a very close relation, see above the use of `[]` for the “indirect” de-referencing of pointers. A field `a[..]` is internally transformed to a pointer on the first

element of the field. Access to the field is then dealt with via the pointer-arithmetics mentioned above. This is also the reason why C/C++-fields always start with 0 (i.e. the localisation of field entries can be done with less operations).

```
int a[20];
int *p = a;           // ok, a is name of the field and can
                      // be used as address to the field
int *p = &(a[0]);    // ... equivalent

p[5] = 4;             // access to a through pointer arithmetic,
                      // a[5] is set to the value of 4
a[5] = 4;             // ... equivalent

int b[20][30];        // b is now an array of 20 arrays of
                      // 30 elements of type int
                      // int *p = b; error! wrong type, since
                      // b is of type pointer to 30 int here
int *r = b[0];        // ok, r points to the first int
                      // in the first (of 20) array(s) of 30 ints
int (*q)[30] = b;     // ok, q points to the first set of 30 int's,
                      // q and b can now be used synonymous
q[12][15] = 26;       // ok, set element (13,16) to the value 26
(*(q+12))[15] = 26;   // ... equivalent
*(*(q+12)+15) = 26;   // ... equivalent
b[12][15] = 26;       // ... equivalent
```

Exercise 7.1.

Practice the use of pointers, using the above examples as a starting point. ■

7.3 Transfer of field-variables to functions

In order to transfer a field to a function, it is sufficient to just transfer the pointer, see sec. 6. For a one-dimensional field this is just the pointer to the first element, while for two-d fields it is a pointer to a field with a compatible, fixed number of elements. Only the first dimension of the field can be left out in this case, all others have to be specified at compile-time. This makes the C-type fields rather in-practical to be used when dynamic memory management is desired.

```
void f( int b[][30] )    // function prototype that transfers b
void g( int (*q)[30] )  // function prototype that transfers

int b[20][30];
int c[10][30];
int d[15][31];

f(b); g(b); f(c); g(c); // ok
f(d); g(d);             // errors, since second dimension not ==30
```

Summary: Transfer of Variables to Functions

As summary and examples for the use of fields in functions, the example in exercise 7 is re-written in order to use functions that use the field:

```
#include <cstdlib>
#include <iostream>
#include <cmath>
using namespace std;

int read_field( int flen, int *fi );
void write_field( int flen, int *fi );
void sort_field( int flen, int *fi );

int main(int argc, char *argv[])
{
    const int Mfmax=20;
    int fi[Mfmax];
    int ifmax=0;
    int ftmp;

    ifmax = read_field( ifmax, fi ); // get field from subroutine
    write_field( ifmax, fi );        // control what you got

    sort_field( ifmax, fi );         // sort the field
    write_field( ifmax, fi );        // control what you got

    system("PAUSE");
    return EXIT_SUCCESS;
}

// function reads the field from another field and returns it back
int read_field( int flen, int *fi )
{
    const int field_length=17;
    int field[field_length] = {1, 0, -10, 12, -2, 5, 5, 1, -4, -2, -4, -8, 3, 25, 51, -4, -5};
    for( int i=0; i<field_length; i++)
    {
        fi[i]=field[i];
    }
    flen = field_length;
    return flen;
}

// function writes field to screen
void write_field( int flen, int *fi )
{
    for(int i=0; i<flen; i++)
```

```
        cout << fi[i] << " ";
    cout << endl;
    return;
}

// function sorts the field and returns it
void sort_field( int flen, int *fi )
{
    int ftmp;
    for(int i=0; i<flen; i++)
    {
        for(int j=flen-1; j>=1; j--)
        {
            if(fi[j]<fi[j-1])
            {
                ftmp=fi[j-1];
                fi[j-1]=fi[j];
                fi[j]=ftmp;
            }
        }
    }
}
```

Exercises

Exercise 7.2.

Consider the following code extract and explain the effect of each line, stating the value

```
int a=15;
int b=10;
```

```
int *p1=&a;
int *p2=&b;
```

```
int **p3=(&p1);
```

```
*p1=5;
*p2=20;
**p3=1;
*p2=*p1;
p1=p2;
*p1=10;
*p2=5;
**p3=8;
```

```
cout << a << "\t" << b << endl;
```



Chapter 8

I/O (Input and Output)

The I/O in C++ is completely renewed as compared to C. This was mostly necessary because of the complex format in C with non-constant number of variables.

The two most important objects are I/O-channels like `istream` and `ostream`. Both are declared in the header `iostream`. In the previous chapters, the standard I/O channels ‘`cin`’ and ‘`cout`’ with the operators ‘`<<`’ or ‘`>>`’ were already used frequently. In the following alternatives and extensions are explained.

In which sequence is I/O processed?

```
#include <iostream>
// The programmer types this
    cout << "Salary: "    << Person.mSalary;

// ... but the compiler reads this:
    ( cout << "Salary: " ) << Person.mSalary;
```

In order to write the string `Salary:` to the screen, the operator `<<` is called from the `ostream` `cout` and returns another `ostream` with which the next `<<` is called. This time, `Person.mSalary` is the second argument. Analogously after an input with `cin`, another `istream` is returned so that one can input again.

8.1 Elementfunctions of `iostreams`

Instead of the form as used up to now, I/O can also be performed in the form of functions – for compatibility. However, the function-form has much wider range of applications, see Tab. 8.1.

```
char c1='x';
cout << c1 << '\n';    // Syntax used up to now
cout.put( c1 );         // Equivalent form using functions
cout.put( '\n' );
```

ostream::put (char c); Example: cout.put(c1);	writes character c to ostream
char c='x'; cout << c;	For comparison ... also writes the char 'c'
istream::get (char c); Example: cin.get(c1);	reads character c from istream
char c; cin >> c;	For comparison ... Reads character c, ignores SPACE, TAB, EOL

Table 8.1: Prototypes of some element-functions for I/O alternative to the use of `cout` und `cin`.

There are many related element-functions like this that allow a much more controlled and well-behaved input and output of data. (Further reading if needed).

8.2 Formatting

Typically it is necessary to change the format of the output, e.g., in order to obtain a better readable table, or in order to prepare the output for the transfer to another software, or to write the full precision on the screen, or ... see Tab. 8.2 for a summary.

Manipulator	Description
int width=12; cout << setw(width);	sets the minimal number of digits for the following output only.
int prec=5; cout << setprecision(prec);	Changes the number of digits behind the comma, or the max. number of digits
char c='*'; cout << setfill(c);	Defines a filling-symbol as, e.g., *.

Table 8.2: Examples of formatting functions for the output

In order to call these functions, first use: `#include<iomanip>`. Note that there are modifications that only act on the next output, and others that remain active permanently.

```
#include<iomanip>
double x=1234.5;
cout << x << '\n';    // leads to: 1234.5
cout << setfill ('*') << setw (10) << setprecision (5);
cout << x << '\n';    // leads to: ****1234.5
cout << x << '\n';    // back to previous setting: 1234.5
```

Furthermore there exist switches (flags) for the formatted output that need no parameters, as summarized in Tab.8.3:

Flag	Group	function
left	adjustfield	left-adjust
right	adjustfield	right-adjust
internal	adjustfield	+/- left
dec	basefield	decimal numbers
hex	basefield	hexadecimal numbers
oct	basefield	octal numbers
showbase		shows the dec-basis of hex- and oct-numbers
showpos		prints '+' before number if positive
uppercase		E, X, A-F instead of: e, x, a-f
fixed	floatfield	Fixed-point notation
scientific	floatfield	scientific notation

Table 8.3: Flags for the formatted output of data

In order to activate flags one can use alternatively

```
cout.setf (ios:: 'flag', ios:: 'group')
cout << setiosflags (ios:: 'flag', ios:: 'group')
```

where the possibilities for 'flag' and 'group' are summarized above in Tab. 8.3, where the second parameter is not always needed. In order to change the setting, or set it to default, for a single flag or for groups of flags, one can use:

```
cout.unsetf (ios:: 'flag')
cout << resetiosflags (ios:: 'flag')

cout.unsetf (ios:: 'group')
cout << resetiosflags (ios:: 'group')
```

8.3 Files

For reading and writing on files, one can use the channels `fstream` and `ofstream`. These are derived from the std-streams `istream` and `ostream`.

```
# include <fstream>
...
double  z=1.234;
int      m=5;
ofstream outs ( "filename", ios::out );
           // open the file with name
```



```

                                // 'filename' for writing
                                // ios::out can be skipped
outs << x << '\n';    // the use of the self-defined output-stream
outs << j << '\n';    // 'outs' is analogous to the use of 'cout'
outs.close();          // close the out-fstream 'outs'

ifstream ins ( "filename", ios::in );
                                // open the file with name
                                // 'filename' to read from
                                // ios::in can be skipped
ins >> z;               // the use of the self-defined input-stream
ins >> m;               // 'ins' is analogous to the use of 'cin'
ins.close();           // close the in-fstream 'ins'

```

At least now, when working with files, we need information whether we have reached the end of the file (EOF) or if the opening of the file was successful. For this, the so-called status-informations can be used, which return – dependent on the status of a file – either `true` or `false`:

```

ins.good()    // true, if no error occurred in stream 'ins'
ins.eof()     // true, if the end of file/stream is reached
ins.bad()     // true, if an hardware error occurred
ins.fail()    // true, if a logical error or a
               // hardware error occurred

ins.clear()   // set ins.good() to true
ins.clear( ios::failbit | ins.rdstate() )
               // set fail() manually to true

```

Example: How to read from a file to its end?

The following examples show which problems can occur doing this.

```

int n;

while (!cin.eof()){
    cin >> n;
    // .... this reads too far
}

while((cin >> n).good){
    // ... this reads not far enough
}

while(cin >> n){
    // ... this is the way it works
}

```

Chapter 9

Some advanced topics

9.1 Dynamical Memory Management

In order to request memory, C++ provides the operators `new` and `new []` for single variables or for fields, respectively. In order to free the memory, `delete` and `delete []` have to be used.

```
int    *p_i  = new int;
int      n = 40;
int    *pa_i = new int[n];
// ...
delete  p_i;
delete[] pa_i;
```

The request for memory requires the type after `new`, and a pointer to this new element is returned. For fields, `new[]` is used and the number of memory-elements within the brackets is requested. The request returns a pointer to the first element in the field.

```
int    *pa_i = new int[n];

pa_i[0]=5;
pa_i[1]=4;
```

The operator `new` returns an “exception” `bad_alloc`, if the request for memory does not succeed and then terminates. It is possible to ‘catch’ error and continue without ending the code, using the commands `try` and `catch` (this will not be covered by this course).

9.2 Object Oriented Programming

9.2.1 Example of Object Oriented Code involving Classes

Just as an example for classes – without further explanations, the following (complete) program. Note that the Class-Definitions appear before the `main()` program, which is especially short.

Class definitions can be placed in separate files as discussed below in the next subsection.

Exercise 9.1.

For Experts: Think about what can be improved in the following code – for this you first will have to understand ... Further reading/learning required. ■

```
#include <stdlib.h>
#include <iostream>
#include <fstream>
using namespace std;
class CSample
{
public:
// destructor
~CSample();

// constructor
CSample(); // ADDED

int Input(const char* fname);
double GetSum();
double GetMax();
int Size;
private:
int Nmax;
double *a;

bool mem_alloc; // ADDED
};
//-----
int CSample::Input(const char* fname)
{
ifstream in(fname);
in >> Nmax;
Size = Nmax;

//memory allocation
a = new double[Nmax]; // ADDED
```

```
mem_alloc=true;

for (int i=0; i<Nmax; i++)
{
in >> a[i];
}
in.close();
return 1;
}
double CSample::GetSum()
{
double sum;
sum = 0;
for (int i=0; i<Nmax; i++)
{
sum = sum + a[i];
}
return sum;
}
double CSample::GetMax()
{
if (a==0) return 0; // ADDED

double m;
m = a[0];
for (int i=1; i<Nmax; i++)
{
if (a[i] > m)
{
m = a[i];
}
}
return m;
}

CSample::CSample() // ADDED
{
mem_alloc = false;
a = 0;
Nmax = 0;
}

CSample::~CSample()
{
// free memory
if (mem_alloc) delete [] a; // ADDED
}

//= MAIN PROGRAM =====
```

```
// Uses the class CSample
// to read in data, print the size of the data-field,
// compute the sum of, and the maximum.
int main()
{
    CSample Test;
    Test.Input("data.dat");
    cout << "Size = " << Test.Size << "\n";
    cout << "Sum = " << Test.GetSum() << "\n";
    cout << "Max = " << Test.GetMax() << "\n";
    cout << endl;
    system("PAUSE");
    return EXIT_SUCCESS;
}
```

9.3 Organisation implementation and header-files

Big programs and program-packages should not be kept in a single file since this can lead to long compilation times and non-clear dealing with and navigation in the file. Furthermore it causes big trouble if more persons work on (a single file) at the same time.

Therefore, one typically splits programs into several files (projects), that can be compiled independently from each other. There are compilable (Implementation-) C++-files with the extension `.cc` or `.cpp`. Files that only contain declarations are called header-files and are used to make functions and data from one file known to another. These have the extension `.h`. In general (except for `main.cc`) an implementation-file comes together with a header-file.

header-files should contain:

1. so-called include-guards, that hinder the multiple inclusion of header-files (required by the standard)

```
#ifndef HEADER_H
#define HEADER_H
    // declarations
#endif
```

2. Declarations of functions that are used amongst several modules. For declarations of local functions, another header-file can be used.
3. Definition of `inline` functions
4. Declaration of classes and their elementfunctionen
5. `extern` declarations of global variables, and declaration of `static` varbiabiles
6. Declaration and definition of `template`-functions and -classes

implementation-files should contain:

1. Implementation of elementfunctions of not-template-classes
2. Implementation of regular functions
3. Declaration von file-scope, class-scope and global static variables

9.3.1 Compiling and linking

Each separate sub-package is compiled to an object file using the `-c` flag e.g. `gcc -c matrix_part.cpp`. This will create an object file called `matrix_part.o`. Then to link the separate object files into an executable call `gcc matrix_part.o other_part.o main_part.o -o my_prog.exe`. The use of a `Makefile` can help automate the process for very large problems (this will not be covered by this course).

Chapter 10

Literature and longer exercises

10.1 Further reading

The many new terms used in the last pages indicates that C++ has much more possibilities, rules and options to be explored. Classes, templates, over-loading, scopes, etc. can not be discussed in this short course. However, here is a list of further literature:

1. Wikipedia ...
2. Brian W. Kernighan and Dennis M. Ritchie, *The C-Programming Language*, 2nd edition, Prentice Hall, 1988.
3. B. Stroustrup, *The C++ Programming Language*, 3rd edition, Addison Wesley, 1997.
4. Steve Oualline, *Practical C++ Programming*, O'Reilly, 1995.
5. C. S. Horstmann, *Mastering C++*, John Wiley & Sons, New York, 1996.
6. S. B. Lippman, *C++ Einführung und Leitfaden*, 3. Auflage, Addison-Wesley, Bonn, 1998.

10.2 Exercises C/C++

Exercise 1

Describe with a few words what the following program is doing. Do not copy this file into your computer - solve this by reading it and write a few sentences (in 15 minutes). Then plot a flow-chart or flow-diagram.

```
#include<iostream>
#include<fstream>
#include<cmath>
using namespace std;

int main(int argc, char *argv[])
{
    // Define field x(t) with length 1000
    double x[1000], t; // output variables
    // initial conditions
    double A, delta;    // A=amplitude, delta=phase-angle
    double mass, ksprng; // mass=mass, ksprng=spring-constant
    double t_max, dt;    // t_max=max-time, dt=time-intervall
    // Request input of the parameters
    cout << "amplitude      "; cin >> A;
    cout << "phase-angle    "; cin >> delta;
    cout << "mass           "; cin >> mass;
    cout << "spring-const. "; cin >> ksprng;
    cout << "max-time       "; cin >> t_max;
    cout << "time-interval "; cin >> dt;
    double omega=sqrt(ksprng/mass);
    // Loop from t=0 to t=t_max
    t=0.0;
    for( int i=0; i<=t_max/dt; i++ )
    {
        x[i]=A*sin(omega*t+delta); // Compute function
        t=t+dt;                    // Step to next time
    }

    ofstream outfile("plot.data"); // Output to file
    t=0.0;
    for( int i=0; i<=t_max/dt; i++ )
    {
        outfile << t << " " << x[i] << "\n";
        t=t+dt;
    }
}
```


Exercise 2

Writing the first program.

Extend the minimal, basic C++ program by the command lines `int inum=0; cin >> inum;` where the second requests input of an integer number from the keyboard.

Write a C++ program that prints the binary representation of the integer number `inum` to the screen. You can check if the program works by, e.g., using the number 496 from the lecture-example.

Exercise 3

What is the greatest value of n that can be used in the sum

$$1^2 + 2^2 + 3^2 + \dots + n^2$$

and gives a total value of the sum less than 100?

(a) Write a short C++ program that answers this question – like in the MATLAB exercise 17.3 Then answer the question also for an exponent of 10 and a total value of the sum of

(b) less than 10^6 and

(c) less than 10^{12} (*Advanced*).

(d) If doing both parts compare the run-time between you C++ and MATLAB implementation.

Exercise 4

Declare an array with 10 001 entries. Compute the solution vector $f(x_i)$, with $x_i = 0, \dots, 10\,000$ (as fast as possible!)

$$f(x_i) = \frac{x_i}{\sin(x_i) + 2}$$

In the same program, form the alternating sum

$$S_n = f(x_0) + f(x_1) - f(x_2) + f(x_3) - \dots$$

up to the last term and print the result to the screen. Finally, compare the result with the result obtained in the MATLAB section 17.3.

Exercise 5

Write a root-finding function in C/C++ and apply to general fifth order equation of your choice. (*Advanced*)

Exercise 6

In the following program (part of which we have discussed today) there are hidden 10 syntax and typo-errors. See how many you can find in 5-10 minutes.

```
#include <iostream>
#include <cmath>
#include <cstdlib>
using namespce std;

int main(int argc, char *argv[])
{
    int i;
    int n=5;
    int *p;
    int a[10];

    for( i=0; i<=10; i++ )
        a[i]=10;
        p[i]=10;

    cout << a << " \n";
    cout << *a << " \n";
    cout << *(a) << " \n";
    cout << a[1] << " \n"
    cout << *(a+2) << " \n";
    cout << size of( int ) << " " << sizeof( int* ) << "\n";
    cout << sizeof( double ) << " " << sizeof( double* ) << "\n";
    cout << " n: " << n << " " << &n << "\n";
    i=n+1;
    cout << " i: " << i << " " << &i << "\n";
    p=&n;
    cout << " p: " << p << " " << *p << "\n";
    *p+=2
    cout << " p: " << p << " " << *p << "\n";
    p++;
    cout << " p: " << p << " " << &p << "\n";

    return 0;
}
```

Exercise 7

This is an example of pointers:

(a) Use the following program segment and play around using the pointers.

```
#include <cstdlib>
#include <iostream>
#include <cmath>
using namespace std;

int main(int argc, char *argv[])
{
    char c='x';
    int a[20];
    int *p = a;           // ok, a is name of the field and can
                        //      be used as address to the field
// int *p = &(a[0]);    // ... equivalent (error->redeclaration)

    p[5] = 4;             // access to a through pointer arithmetic,
                        // a[5] is set to the value of 4
    a[5] = 4;             // ... equivalent
                        //
    int b[20][30];        // b is now an array of 20 arrays of
                        //      30 elements of type int
                        //      int *p = b; error! wrong type, since
                        //      b is of type pointer to 30 int here
    int *r = b[0];        // ok, r points to the first int
                        //      in the first (of 20) array(s) of 30 ints
    int (*q)[30] = b;     // ok, q points to the first set of 30 int's,
                        //      q and b can now be used synonymous
    q[12][15] = 26;       // ok, set element (13,16) to the value 26
    (*(q+12))[15] = 26;   // ... equivalent
    (*(q+12)+15) = 26;    // ... equivalent
    b[12][15] = 26;       // ... equivalent

    cout.put(c);
    int *p_i = new int;
//    Person *p_p = new Person("B.Stroustrup",45);
    int n = 40;
    int *pa_i = new int[n];
// ...
    delete p_i;
//    delete p_p;
    delete[] pa_i;
//
//
    pa_i[0]=5;
```

```
    pa_i[1]=4;

    for(int i=0; i<10; i++) cout << pa_i[i] << ' ';

    return 0;
}
```

Serious programming exercise:

Define an `int` field (vector) that contains the numbers:

```
{ 1, 0, -10, 12, -2, 5, 5, 1, -4, -2, -4, -8, 3, 25, 51, -4, -5 }
```

Then sort the entries of this linear field/vector.

- (b) Read the field from a file.
- (c) Use a linked-list to store the field.
- (d) Implement a linked list (with pointers) to store the field.
- (e) Use this linked list (with pointers) to sort the field.
- (f) Write the sorted list to an output file, but remove double entries from the list. What is the length of the field before and after this sorting/remove operation?
- (g) (Advanced) Implement a (heap) tree structure to sort the field such that the largest entry is always on top of the tree.

Part: B

MATLAB

Chapter 11

Getting started with MATLAB

MATLAB is a tool for mathematical (technical) calculations. Firstly, it can be used as a scientific calculator. Next, it allows you to plot or visualize data in many different ways, perform matrix algebra, work with polynomials or integrate functions. Like in a programmable calculator, you can create, execute and save a sequence of commands in order to make your computational process automatic. It can be used to store or retrieve data. In the end, MATLAB can also be treated as a user-friendly programming language, which gives the possibility to handle mathematical calculations in an easy way. In summary, as a computing/programming environment, MATLAB is especially designed to work with data sets as a whole such as vectors, matrices and images. Therefore, PRTOOLS (<http://prtools.org>), a toolbox for Pattern Recognition, and DIPIMAGE (<http://www.ph.tn.tudelft.nl/DIPlib/dipimage.1.html>), a toolbox for Image Processing, have been developed under MATLAB.

Under Windows, you can start MATLAB by double clicking on the MATLAB icon that should be on the desktop of your computer. On a unix system, type `matlab` at the command line. Running MATLAB creates one or more windows on your screen. The most important is the *Command Window*, which is the place where you interact with MATLAB, i.e. it is used to enter commands and display text results. The string `>>` is the MATLAB prompt (or `EDU>>` for the Student Edition). When the Command Window is active, a cursor appears after the prompt, indicating that MATLAB is waiting for your command. MATLAB responds by printing text in the Command Window or by creating a *Figure Window* for graphics. To exit MATLAB use the command `exit` or `quit`. `Control-C` is a local abort which kills the current execution of a command.

11.1 Input via the command-line

MATLAB is an interactive system; commands followed by **Enter** are executed immediately. The results are, if desired, displayed on screen. However, execution of a command will be possible if the command is typed according to the rules. Table 11.1 shows a list of commands used to solve

indicated mathematical equations (a , b , x and y are numbers). Below you find basic information to help you starting with MATLAB :

- Commands in MATLAB are executed by pressing **Enter** or **Return**. The output will be displayed on screen immediately. Try the following (hit **Enter** after the end of line):

```
1 >> 3 + 7.5
2 >> 18/4
3 >> 3 * 7
```

Note that spaces are **not** important in MATLAB.

- The result of the last performed computation is ascribed to the variable **ans**, which is an example of a MATLAB built-in variable. It can be used in the subsequent command. For instance:

```
1 >> 14/4
2 ans =
3      3.5000
4 >> ans^(-6)
5 ans =
6      5.4399e-04
```

$5.4399\text{e-}04$ is a computer notation of $5.4399 * 10^{-4}$ (see **Preliminaries**). Note that **ans** is always overwritten by the last command that has no assignment.

- You can also define your own variables. Look how the information is stored in the variables **a** and **b**:

```
1 >> a = 14/4
2 a =
3      3.5000
4 >> b = a^(-6)
5 b =
6      5.4399e-04
```

Read **Preliminaries** to better understand the concept of variables. You will learn more on MATLAB variables in section **12.1.2**.

I can not find the Preliminaries section

- When the command is followed by a semicolon ';', the output is suppressed. Check the difference between the following expressions:

```
1 >> 3 + 7.5
2 >> 3 + 7.5;
```

- It is possible to execute more than one command at the same time; the separate commands should then be divided by commas (to display the output) or by semicolons (to suppress the output display), e.g.:

```
1 >> sin(pi/4), cos(pi); sin(0)
2 ans =
3      0.7071
4 ans =
5      0
```

Note that the value of **cos(pi)** is not printed.

- By default, MATLAB displays only 5 digits. The command `format long` increases this number to 15, `format short` reduces it to 5 again. For instance:

```

1 >> 312/56
2 ans =
3     5.5714
4 >> format long
5 >> 312/56
6 ans =
7     5.57142857142857

```

- The output may contain some empty lines; this can be suppressed by the command `format compact`. In contrast, the command `format loose` will insert extra empty lines.
- To enter a statement that is too long to be typed in one line, use three periods `'...'` followed by **Enter** or **Return**. For instance:

```

1 >> sin(1) + sin(2) - sin(3) + sin(4) - sin(5) + sin(6) - ...
2     sin(8) + sin(9) - sin(10) + sin(11) - sin(12)
3 ans =
4     1.0357

```

- MATLAB is case sensitive, for example, *a* is written as **a** in MATLAB; **A** will result in an error in this case.
- All text after a percent sign `%` until the end of a line is treated as a comment. Enter e.g. the following:

```

1 >> sin(3.14159)           % this is an approximation of sin(pi)

```

You will notice that some examples in this text are followed by comments. They are meant for you and you should skip them while typing.

- Previous commands can be fetched back with the `↑`-key. The command can also be changed, the `←` and `→`-keys may be used to move around in a line and edit it. In case of a long line, `Ctrl-a` and `Ctrl-e` might be useful; they allow to move the cursor at the beginning or the end of the line, respectively.
- To recall the most recent command starting from e.g. `c`, type `c` at the prompt followed by the `↑`-key. Similarly, `cos` followed by the `↑`-key will find the last command starting from `cos`.

Since MATLAB executes the command immediately, it might be useful to have an idea of the expected outcome. You might be surprised how long it takes to print out a 1000×1000 matrix!

Mathematical notation	MATLAB command
$a + b$	<code>a + b</code>
$a - b$	<code>a - b</code>
ab	<code>a * b</code>
$\frac{a}{b}$	<code>a / b</code> or <code>b \ a</code>
x^b	<code>x^b</code>
\sqrt{x}	<code>sqrt(x)</code> or <code>x^0.5</code>
$ x $	<code>abs(x)</code>
π	<code>pi</code>
$4 \cdot 10^3$	<code>4e3</code> or <code>4*10^3</code>
$i(\sqrt{-1})$	<code>i</code> or <code>j</code>
$3 - 4i$	<code>3-4*i</code> or <code>3-4*j</code>
e, e^x	<code>exp(1)</code> , <code>exp(x)</code>
$\ln x, \log x$	<code>log(x)</code> , <code>log10(x)</code>
$\sin x, \arctan x, \dots$	<code>sin(x)</code> , <code>atan(x)</code> ,...

Table 11.1: Mathematical notation in MATLAB 's commands.

11.2 help-facilities

There are several ways for getting help with MATLAB. A good way to start is to type:

```
1 >> helpdesk
```

The helpdesk command opens an HTML document, which contains the MATLAB reference book. This is quite a large document! Useful sections for beginners include Getting Started and MATLAB Functions. You might want to browse them before beginning your first session. Its helpful to keep this browser window open while you are working with MATLAB , so that you can refer to it easily.

There are other ways of getting help with MATLAB that do not make you search through the whole reference book. If you already know the name of a MATLAB function, for example, `quit`, and you want to learn about its use, enter:

```
1 >> help quit
```

You will see a description of the command `quit`. During your first experiences with MATLAB you can from time to time take a:

```
1 >> demo
```

session which will show you various features of MATLAB . Another very useful feature is the `lookfor` command. It looks for all the commands related to a given topic. Try:

```
1 >> lookfor 'help'
```

It will list all of the commands we just discussed.

Besides the inherent help of MATLAB, you will find various online courses and references in the internet and, last but not least, textbooks. Here, we just mention a view:

Online sources

- MATLAB home page (from manufacturers)
<http://www.mathworks.com>
- MATLAB online course (TU Eindhoven)
<http://www.imc.tue.nl>
- MATLAB online Reference Documentation
<http://www.math.ufl.edu/help/matlab/ReferenceTOC.html>

Books and others references

- MATLAB An Introduction with Applications, Amos Gilat, 2008.
- MATLAB Programming for Engineers, Stephen Chapman, 2007.
- Essential MATLAB for Engineers and Scientists, Brian Hahn and Dan Valentine, 2007.

11.3 Interrupting a command or program

Sometimes you might spot an error in your command or program. Due to this error it can happen that the command or program does not stop. Pressing **Ctrl-C** (or **Ctrl-Break** on PC) forces MATLAB to stop the process. Sometimes, however, you may need to press a few times. After this the MATLAB prompt (**>>**) re-appears. This may take a while, though.


11.4 Path

In MATLAB, commands or programs are contained in m-files, which are just plain text files and have an extension `'.m'`. The m-file must be located in one of the directories which MATLAB automatically searches. The list of these directories can be listed by the command `path`. One of the directories that is always taken into account is the *current working directory*, which can be identified by the command `pwd`. Use `path`, `addpath` and `rmpath` functions to modify the path. It is also possible to access the path browser from the **File** menu-bar, instead.

Exercise 11.1.

Type `path` to check which directories are placed on your path. Add your personal directory to the path (assuming that you created your personal directory for working with MATLAB). ■

11.5 Workspace issues

If you work in the Command Window, MATLAB memorizes all commands that you entered and all variables that you created. These commands and variables are said to reside in the MATLAB *workspace*. They might be easily recalled when needed, e.g. to recall previous commands, the -key is used. Variables can be verified with the commands `who`, which gives a list of variables present in the workspace, and `whos`, which includes also information on name, number of allocated bytes and class of variables. For example, assuming that you performed all commands from section 1.1, after typing `who` you should get the following information:

```
1 >> who
2 Your variables are:
3 a          ans          b          x
```

The command `clear <name>` deletes the variable `<name>` from the MATLAB workspace, `clear` or `clear all` removes all variables. This is useful when starting a new exercise. For example:

```
1 >> clear a x
2 >> who
3 Your variables are:
4 ans          b
```

Note that you cannot use comma after a variable, i.e. `clear a, x`, as it will be interpreted in MATLAB as `clear a` and print `x` on the screen. See what is the result of:

```
1 >> clear all
2 >> a = 1; b = 2; c = 3;
3 >> clear a, b, c
```

11.6 Saving and loading data

The easiest way to save or load MATLAB variables is by using (clicking) the **File** menu-bar, and then selecting the **Save Workspace as...** or **Load Workspace...** items respectively. Also MATLAB commands exist which save data to files and which load data from files.

The command `save` allows for saving your workspace variables either into a binary file or an ASCII file (check **Preliminaries** on binary and ASCII files). Binary files automatically get the `'.mat'` extension, which is not true for ASCII files. However, it is recommended to add a `'.txt'` or `.dat` extension.

Exercise 11.2.

Learn how to use the `save` command by exercising:

```
1 >> clear all
2 >> s1 = sin(pi/4);
3 >> c1 = cos(pi/4); c2 = cos(pi/2);
4 >> str = 'hello world';           % this is a string
```

```

5 >> save                                % saves all variables in binary format to
6                                         % matlab.mat
7 >> save data                            % saves all variables in binary format to data.mat
8 >> save numdata s1, c1                  % saves numeric variables s1 and c1 to numdata.mat
9 >> save strdata str                     % saves a string variable str to strdata.mat
10 >> save allcos.dat c* -ascii           % saves c1,c2 in 8-digit ascii format to
11                                         % allcos.dat

```

■

The `load` command allows for loading variables into the workspace. It uses the same syntax as `save`.

Exercise 11.3.

Assuming that you have done the previous exercise, try to load variables from the created files. Before each `load` command, clear the workspace and after loading check which variables are present in the workspace (use `who`).

```

1 >> load                                % loads all variables from the file matlab.mat
2 >> load data s1 c1                     % loads only specified variables from the file
3                                         % data.mat
4 >> load strdata                         % loads all variables from the file strdata.mat

```

It is also possible to read ASCII files that contain rows of space separated values. Such a file may contain comments that begin with a percent character. The resulting data is placed into a variable with **the same** name as the ASCII file (without the extension). Check, for example:

```

1 >> load allcos.dat                     % loads data from allcos.dat into variable allcos
2 >> who                                 % lists variables present in the workspace now

```

■

11.7 Exercises

Exercise 11.4.

- Is the inverse cosine function, known as \cos^{-1} or *arccos*, one of the MATLAB's elementary functions?
- Does MATLAB have a mathematical function to calculate the greatest common divisor?
- Look for information on logarithms.

Use `help` or `lookfor` to find out. ■

Chapter 12

Basic syntax and variables

12.1 MATLAB as a calculator

There are three kinds of numbers used in MATLAB: integers, real numbers and complex numbers. In addition, MATLAB has representations of the non-numbers: `Inf`, for positive infinity, generated e.g. by `1/0`, and `NaN`, Not-a-Number, obtained as a result of the mathematically undefined operations such as `0/0` or $\infty - \infty$.

You have already got some experience with MATLAB and you know that it can be used as a calculator. To do that you can, for example, simply type:

```
1 >> (23*17)/7
```

The result will be:

```
1 ans =  
2     55.8571
```

MATLAB has six basic arithmetic operations, such as: `+`, `-`, `*`, `/` or `\` (right and left divisions) and `^` (power). Note that the two division operators are different:

```
1 >> 19/3                                % mathematically: 19/3  
2 ans =  
3     6.3333  
4 >> 19\3, 3/19                          % mathematically: 3/19  
5 ans =  
6     0.1579  
7 ans =  
8     0.1579
```

Basic built-in functions, trigonometric, exponential, etc, are available for a user. Try `help elfun` to get the list of elementary functions.

Exercise 12.1.

Evaluate the following expressions by hand and use MATLAB to check the answers. Note the difference between the left and right divisors. Use `help` to learn more on commands rounding numbers, such as: `round`, `floor`, `ceil`, etc.

- | | |
|---------------------------------|--------------------------------------------------------------|
| • $2/2 * 3$ | • $3^2/4$ |
| • $8 * 5 \backslash 4$ | • $3^2 \wedge 3$ |
| • $8 * (5 \backslash 4)$ | • $2 + \text{round}(6/9 + 3 * 2)/2$ |
| • $7 - 5 * 4 \backslash 9$ | • $2 + \text{floor}(6/9 + 3 * 2)/2$ |
| • $6 - 2/5 + 7^2 - 1$ | • $2 + \text{ceil}(6/9 + 3 * 2)/2$ |
| • | • |
| $10/2 \backslash 5 - 3 + 2 * 4$ | $x = \text{pi}/3, x = x - 1, x = x + 5, x = \text{abs}(x)/x$ |

■

INTERMEZZO

12.1.1 An introduction to floating-point numbers

In a computer, numbers can be represented only in a *discrete* form. It means that numbers are stored within a limited range and with a finite precision. Integers can be represented *exactly* with the base of 2 (read **Preliminaries** on bits and the binary system). The typical size of an integer is 16 bits, so the largest positive integer, which can be stored, is $2^{16} = 65536$. If negative integers are permitted, then 16 bits allow for representing integers between -32768 and 32767 . Within this range, operations defined on the set of integers can be performed exactly.

However, this is not valid for other real numbers. In practice, computers are integer machines and are capable of representing real numbers only by using complicated codes. The most popular code is the *floating point* standard. The term floating point is derived from the fact that there is no fixed number of digits before and after the decimal point, meaning that the decimal point can float. Note that most floating-point numbers that a computer can represent are just approximations. Therefore, care should be taken that these approximations lead to reasonable results. If a programmer is not careful, small discrepancies in the approximations can cause meaningless results. Note the difference between e.g. the integer arithmetic and floating-point arithmetic:

Integer arithmetic:

$$2 + 4 = 6$$

$$3 * 4 = 12$$

$$25/11 = 2$$

Floating-point arithmetic

$$18/7 = 2.5714$$

$$2.5714 * 7 = 17.9998$$

$$10000/3 = 3.3333\text{e}+03$$

When describing floating-point numbers, precision refers to the number of bits used for the fractional part. The larger the precision, the more exact fractional quantities can be represented. Floating-point numbers are often classified as single precision or double precision. A double-precision number uses twice as many bits as a single-precision value, so it can represent fractional values much better. However, the precision itself is not double. The extra bits are also used to increase the range of magnitudes that can be represented.

MATLAB relies on a computer's floating point arithmetic. You could have noticed that in the last exercise since the value of $\sin(\pi)$ was almost zero, and not completely zero. It came from the fact that both the value of π is represented with a finite precision and the sin function is also approximated.

The fundamental type in MATLAB is `double`, which stands for a representation with a double precision. It uses 64 bits. The single precision obtained by using the `single` type offers 32 bits. Since most numeric operations require high accuracy the `double` type is used by default. This means, that when the user is inputting integer values in MATLAB (for instance, `k = 4`), the data is still stored in `double` format.

The relative accuracy might be defined as the smallest positive number ϵ that added to 1, creates the result larger than 1, i.e. $1 + \epsilon > 1$. It means that in floating-point arithmetic, for positive values smaller than ϵ , the result equals to 1 (in exact arithmetic, of course, the result is always larger than 1). In MATLAB, ϵ is stored in the built-in variable `eps` $\approx 2.2204e-16$. This means that the relative accuracy of individual arithmetic operations is about 15 digits.

END INTERMEZZO

12.1.2 Assignments and variables

Working with complex numbers is easily done with MATLAB.

Exercise 12.2.

Choose two complex numbers, for example $-3 + 2i$ and $5 - 7i$. Add, subtract, multiply, and divide these two numbers. ■

During this exercise, the complex numbers had to be typed four times. To reduce this, assign each number to a variable. For the previous exercise, this results in:

```
1 >> z = -3 + 2*i; w = 5 - 7*i;
2 >> y1 = z + w; y2 = z - w;
3 >> y3 = z * w;
4 >> y4 = z / w; y5 = w \ z;
```

Formally, there is no need to declare (i.e. define the name, size and the type of) a new variable in MATLAB. A variable is simply created by an assignment (e.g. `z = -3 + 2*i`), i.e. values are assigned to variables. Each newly created numerical variable is *always* of the `double` type, i.e. real numbers are approximated with the highest possible precision. You can change this type by converting it into e.g. the `single` type¹. In some cases, when huge matrices should be handled and precision is not very important, this might be a way to proceed. Also, when only integers are taken into consideration, it might be useful to convert the `double` representations into e.g. `int32`¹ integer type. Note that integer numbers are represented exactly, no matter which numeric type is used, as long as the number can be represented in the number of bits used in the numeric type.

Bear in mind that *undefined* values cannot be assigned to variables. So, the following is not possible:

¹a variable `a` is converted into a different type by performing e.g. `a = single(a)`, `a = int32(a)` etc.

```
1 >> clear x; % to make sure that x does not exist
2 >> f = x^2 + 4 * sin(x)
```

It becomes possible by:

```
1 >> x = pi / 3; f = x^2 + 4 * sin(x)
```

Variable name begins with a letter, followed by letters, numbers or underscores. MATLAB recognizes only first 31 characters of the name.

Exercise 12.3.

Here are some examples of different types of MATLAB variables. You do not need to understand them all now, since you will learn more about them during the course. Create them manually in MATLAB:

```
1 >> this_is_my_very_simple_variable_today = 5 % check what happens; the name is
2 % very long
3 >> 2t = 8 % what is the problem with this
4 % command?
5 >> M = [1 2; 3 4; 5 6] % a matrix
6 >> c = 'E' % a character
7 >> str = 'Hello world' % a string
8 >> m = ['J', 'o', 'h', 'n'] % try to guess what it is
```

Check the types by using the command `whos`. Use `clear <name>` to remove a variable from the workspace. ■

As you already know, MATLAB variables can be created by an assignment. There is also a number of built-in variables, e.g. `pi`, `eps` or `i`, summarized in Table 12.1. In addition to creating variables by assigning values to them, another possibility is to copy one variable, e.g. `b` into another, e.g. `a`. In this way, the variable `a` is automatically created (if `a` already existed, its previous value is lost):

```
1 >> b = 10.5;
2 >> a = b;
```

A variable can also be created as a result of the evaluated expression:

```
1 >> a = 10.5; c = a^2 + sin(pi*a)/4;
```

or by loading data from text or `*.mat` files.

If `min` is the name of a function (see `help min`), then `a` defined, e.g. as:

```
1 >> b = 5; c = 7;
2 >> a = min(b,c); % create a as the minimum of b and c
```

will call that function, with the values `b` and `c` as parameters. The result of this function (its return value) will be written (assigned) into `a`. So, variables can be created as results of the execution of built-in or user-defined functions (you will learn more how to built own functions in section 17.1).

Important: do not use variable names which are defined as function names (for instance `mean` or `error`)², because from this point you cannot access the function. For example try

```

1 >> mean([4 6])
2
3 ans =
4
5      5
6
7 >> mean=3
8
9 mean =
10
11      3
12
13 >> mean([4 6])
14 ??? Index exceeds matrix dimensions.
```

At the second calling `mean` is a variable not the function `mean`.

If you are going to use a suspicious variable name, use `help <name>` to find out if the function already exists.

Variable name	Value/meaning
<code>ans</code>	the default variable name used for storing the last result
<code>pi</code>	$\pi = 3.14159\dots$
<code>eps</code>	the smallest positive number that added to 1 makes a result larger than 1
<code>inf</code>	representation for positive infinity, e.g. $1/0$
<code>nan</code> or <code>NaN</code>	representation for not-a-number, e.g. $0/0$
<code>i</code> or <code>j</code>	$i = j = \sqrt{-1}$
<code>nargin/nargout</code>	number of function input/output arguments used
<code>realmin/realmax</code>	the smallest/largest usable positive real number: $1.7977e+308$ / $2.2251e-308$

Table 12.1: Built-in variables in MATLAB.

12.2 Exercises

Exercise 12.4.

Define the format in MATLAB such that empty lines are suppressed and the output is given with 15 digits. Calculate:

```

1 >> pi
2 >> sin(pi)
```

Note that the answer is not exactly 0. Use the command `format` to put MATLAB in its standard-format. ■

²There is always one exception of the rule: variable `i` is often used as counter in a loop, while it is also used as $i = \sqrt{-1}$.

Chapter 13

Mathematics with vectors and matrices

The basic element of MATLAB is a matrix (or an array). Special cases are:

- a 1×1 -matrix: a scalar or a single number;
- a matrix existing only of one row or one column: a vector. [Actually this is stored as a two matrix with one dimension being length 1].

Note that MATLAB may behave differently depending on the input, whether it is a number, a vector or a two-dimensional (or more-dimensional) matrix.

13.1 Vectors

Row vectors are lists of numbers separated either by commas or by spaces. They are examples of simple *arrays*. First element has index 1. The number of entries is known as the *length* of the vector (the command `length` exists as well). Their entities are referred to as *elements* or *components*. The entries must be enclosed in `[]`:

```
1 >> v = [-1 sin(3) 7]
2 v =
3     -1.0000     0.1411     7.0000
4 >> length(v)
5 ans =
6      3
```

A number of operations can be done on vectors. A vector can be multiplied by a scalar, or added/subtracted to/from another vector with *the same* length, or a number can be added/-subtracted to/from a vector. All these operations are carried out element-by-element. Vectors can also be built from the already existing ones.

```

1 >> v = [-1 2 7]; w = [2 3 4];
2 >> z = v + w                                % an element-by-element sum
3 z =
4      1      5     11
5 >> vv = v + 2                                % add 2 to all elements of vector v
6 vv =
7      1      4      9
8 >> t = [2*v, -w]
9 ans =
10     -2      4     14     -2     -3     -4

```

Also, a particular value can be changed or displayed:

```

1 >> v(2) = -1                                % change the 2nd element of v
2 v =
3     -1     -1      7
4 >> w(2)                                    % display the 2nd element of w
5 ans =
6      3

```

13.1.1 Colon notation

A colon notation is an important shortcut, used when producing row vectors (see Table 13.1 and `help colon`):

```

1 >> 2:5
2 ans =
3      2      3      4      5
4 >> -2:3
5 ans =
6     -2     -1      0      1      2      3

```

In general, `first:step:last` produces a vector of entities with the value `first`, incrementing by the `step` until it reaches `last`:

```

1 >> 0.2:0.5:2.4
2 ans =
3      0.2000      0.7000      1.2000      1.7000      2.2000
4 >> -3:3:10
5 ans =
6     -3      0      3      6      9
7 >> 1.5:-0.5:-0.5 % negative step is also possible
8 ans =
9      1.5000      1.0000      0.5000      0     -0.5000

```

13.1.2 Extracting and appending parts of a vector

Parts of vectors can be extracted by using a colon notation:

```

1 >> r = [-1:2:6, 2, 3, -2]      % -1:2:6 => -1  1  3  5
2 r =
3     -1     1     3     5     2     3    -2
4 >> r(3:6)                      % get elements of r which are on the positions
5                                % from 3 to 6
6 ans =
7     3     5     2     3
8 >> r(1:2:5)                    % get elements of r which are on the positions
9                                % 1, 3 and 5
10 ans =
11    -1     3     2
12 >> r(5:-1:2)                  % what will you get here?

```

MATLAB allocates memory for all variables on the fly. This allows you to increase the size of a vector simply by assigning a value to an element that has not been previously used.

```

1 >> x = linspace(21,25,5)
2
3 x =
4
5     21     22     23     24     25
6
7 >> x(7) = -9
8
9 x =
10
11     21     22     23     24     25     0    -9

```

13.1.3 Column vectors and transposing

To create *column* vectors, you should separate entries by new lines or by a semicolon ';':

```

1 >> z = [1
2         7
3         7];
4 z =
5     1
6     7
7     7
8 >> u = [-1; 3; 5]
9 u =
10    -1
11     3
12     5

```

The operations applied to row vectors can be applied to column vectors, as well. You cannot, however, add a column vector to a row vector. To do that, you need an operation called *transposing*, which converts a column vector into a row vector and vice versa:

```

1 >> u'                          % u is a column vector and u' is a row vector
2 ans =
3    -1     3     5

```

```

4 >> v = [-1 2 7];           % v is a row vector
5 >> u + v                     % you cannot add a column vector u to a row
6                               % vector v
7 ??? Error using ==> +
8 Matrix dimensions must agree.
9 >> u' + v
10 ans =
11     -2     5    12
12 >> u + v'
13 ans =
14     -2
15     5
16     12

```

If \mathbf{z} is a complex vector, then \mathbf{z}' gives the conjugate transpose of \mathbf{z} . For instance:

```

1 >> z = [1+2i, -1+i]
2 z =
3     1.0000 + 2.0000i    -1.0000 + 1.0000i
4 >> z'                               % this is the conjugate transpose
5 ans =
6     1.0000 - 2.0000i
7    -1.0000 - 1.0000i
8 >> z.'                               % this is the traditional transpose
9 ans =
10     1.0000 + 2.0000i
11    -1.0000 + 1.0000i

```

13.1.4 Product, divisions and powers of vectors

You can now compute the inner product between two vectors \mathbf{x} and \mathbf{y} of the same length, $\mathbf{x}^T \mathbf{y} = \sum_i x_i y_i$, in a simple way:

```

1 >> u = [-1; 3; 5]           % a column vector
2 >> v = [-1; 2; 7]           % a column vector
3 >> u * v                     % you cannot multiply a column vector by a column
4                               % vector
5 ??? Error using ==> *
6 Inner matrix dimensions must agree.
7 >> u' * v                     % this is the inner product
8 ans =
9     42

```

Another way to compute the inner product is by the use of the dot product, i.e. `.*`, which performs element-wise multiplication. Given two vectors \mathbf{x} and \mathbf{y} of the same length, an element-wise multiplication is defined as a vector $[x_1 y_1, x_2 y_2, \dots, x_n y_n]$, thus, the corresponding elements of two vectors are multiplied. For instance:

```

1 >> u .* v                     % this is an element-by-element multiplication
2     1
3     6
4    35
5 >> sum(u.*v)                  % this is an another way to compute the inner product
6 ans =

```

```

7      42
8  >> z = [4 3 1];           % z is a row vector
9  >> sum(u'.*z)              % this is the inner product
10 ans =
11     10
12 >> u'*z'                   % since z is a row vector, u'*z' is the inner product
13 ans =
14     10

```

You can now easily tabulate the values of a function for a given list of arguments. For instance:

```

1 >> x = 1:0.5:4;
2 >> y = sqrt(x) .* cos(x)
3 y =
4     0.5403     0.0866    -0.5885    -1.2667    -1.7147    -1.7520    -1.3073

```

Mathematically, a division of one vector by another is an undefined operation. However, in MATLAB, the operator `./` is introduced to perform an element-by-element division. It is, therefore, defined for vectors of the same size and type:

```

1 >> x = 2:2:10
2 x =
3     2     4     6     8    10
4 >> y = 6:10
5 y =
6     6     7     8     9    10
7 >> x./y
8 ans =
9     0.3333     0.5714     0.7500     0.8889     1.0000
10 >> z = -1:3
11 z =
12    -1     0     1     2     3
13 >> x./z           % division 4/0, resulting in Inf
14 Warning: Divide by zero.
15 ans =
16    -2.0000         Inf     6.0000     4.0000     3.3333
17 >> z./z           % division 0/0, resulting in NaN
18 Warning: Divide by zero.
19 ans =
20     1    NaN     1     1     1

```

The operator `./` can also be used to divide a scalar by a vector:

```

1 >> x=1:5; 2/x           % this is not possible
2 ??? Error using ==> /
3 Matrix dimensions must agree.
4 >> 2./x                 % but this is!
5 ans =
6     2.0000     1.0000     0.6667     0.5000     0.4000

```

Exercise 13.1.

Get acquainted with operations on row and column vectors. Perform, for instance:

- Create a vector consisting of the even numbers between 21 and 47.

- Let $\mathbf{x} = [4 \ 5 \ 9 \ 6]$.
 - Subtract 3 from each element.
 - Add 11 to the odd-index elements.
 - Compute the square root of each element.
 - Raise to the power 3 each element.
- Create a vector \mathbf{x} with the elements:
 - 2, 4, 6, 8, ..., 16
 - 9, 7, 5, 3, 1, -1, -3, -5
- Given $\mathbf{x} = [2 \ 1 \ 3 \ 7 \ 9 \ 4 \ 6]$, explain what the following commands do (note that $\mathbf{x}(\text{end})$ points to the last element of \mathbf{x}):
 - $\mathbf{x}(3)$
 - $\mathbf{x}(1:7)$
 - $\mathbf{x}(1:\text{end})$
 - $\mathbf{x}(1:\text{end}-1)$
 - $\mathbf{x}(2:2:6)$
 - $\mathbf{x}(6:-2:1)$
 - $\mathbf{x}(\text{end}-2:-3:2)$
 - $\text{sum}(\mathbf{x})$
 - $\text{mean}(\mathbf{x})$
 - $\text{min}(\mathbf{x})$

■

Command	Result
$A(i, j)$	A_{ij}
$A(:, j)$	j -th column of A
$A(i, :)$	i -th row of A
$A(k:l, m:n)$	$(l - k + 1) \times (n - m + 1)$ matrix with elements A_{ij} with $k \leq i \leq l, m \leq j \leq n$
$v(i:j)'$	'vector-part' $(v_i, v_{i+1}, \dots, v_j)$ of vector v

Table 13.1: Manipulation of (groups of) matrix elements.

13.2 Matrices

Row and column vectors are special types of matrices. An $n \times k$ matrix is a rectangular array of numbers having n rows and k columns. Defining a matrix in MATLAB is similar to defining a vector. The generalization is straightforward, if you see that a matrix consists of row vectors (or column vectors). Commas or spaces are used to separate elements in a row, and semicolons

are used to separate individual rows. For example, the matrix $A = \begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix}$ is defined as:

```

1 >> A = [1 2 3; 4 5 6; 7 8 9]    % row by row input
2 A =
3     1     2     3
4     4     5     6
5     7     8     9

```

Command	Result
<code>n = rank(A)</code>	n becomes the rank of matrix A
<code>x = det(A)</code>	x becomes the determinant of matrix A
<code>x = size(A)</code>	x becomes a row-vector with 2 elements: the number of rows and columns of A
<code>x = trace(A)</code>	x becomes the trace (sum of diagonal elements) of matrix A
<code>x = norm(v)</code>	x becomes the Euclidean length of vector v
<code>C = A + B</code>	sum of two matrices
<code>C = A - B</code>	subtraction of two matrices
<code>C = A * B</code>	multiplication of two matrices
<code>C = A .* B</code>	'element-by-element' multiplication (A and B are of equal size)
<code>C = A^k</code>	power of a matrix ($k \in \mathbb{Z}$; can also be used for A^{-1})
<code>C = A.^k</code>	'element-by-element' power of a matrix
<code>C = A'</code>	the transposed of a matrix; A^T
<code>C = A ./ B</code>	'element-by-element' division (A and B are of equal size)
<code>X = A \ B</code>	finds the solution in the least squares sense to the system of equations $AX = B$
<code>X = B / A</code>	finds the solution of $XA = B$, analogous to the previous command
<code>C = inv(A)</code>	C becomes the inverse of A
<code>C = null(A)</code>	C is an orthonormal basis for the null space of A obtained from the singular value decomposition
<code>C = orth(A)</code>	C is an orthonormal basis for the range of A
<code>L = eig(A)</code>	L is a vector containing the (possibly complex) eigenvalues of a square matrix A
<code>[Q,L] = eig(A)</code>	produces a diagonal matrix L of eigenvalues and a full matrix Q whose columns are the corresponding eigenvectors of a square matrix A
<code>S = svd(A)</code>	S is a vector containing the singular values of a rectangular matrix A
<code>[U,S,V] = svd(A)</code>	S is a diagonal matrix with nonnegative diagonal elements in decreasing order; columns of U and V are the accompanying singular vectors of A
<code>x = linspace(a,b,n)</code>	generates a vector x of n equally spaced points between a and b
<code>x = logspace(a,b,n)</code>	generates a vector x starting at 10^a and ended at 10^b containing n values
<code>A = eye(n)</code>	A is an $n \times n$ identity matrix
<code>A = zeros(n,m)</code>	A is an $n \times m$ matrix with zeros (default $m = n$)
<code>A = ones(n,m)</code>	A is an $n \times m$ matrix with ones (default $m = n$)
<code>A = diag(v)</code>	results in a diagonal matrix with the elements v_1, v_2, \dots, v_n on the diagonal
<code>v = diag(A)</code>	results in a vector equivalent to the diagonal of A
<code>X = tril(A)</code>	X is lower triangular part of A
<code>X = triu(A)</code>	X is upper triangular part of A
<code>A = rand(n,m)</code>	A is an $n \times m$ matrix of elements drawn from a uniform distribution on $[0, 1]$
<code>A = randn(n,m)</code>	A is an $n \times m$ matrix of elements drawn from a standard normal distribution
<code>v = max(A)</code>	v is a vector of the maximum values of the columns in A
<code>v = max(A,dim)</code>	v is a vector of the maximum values along the dimension dim in A
<code>v = min(A),</code> <code>v = min(A,dim)</code>	ditto - with minimum
<code>v = sum(A),</code> <code>v = sum(A,dim)</code>	ditto - with sum

Table 13.2: Frequently used matrix operations and functions.

Other examples are, for instance:

```

1 >> A2 = [1:4; -1:2:5]
2 A2 =
3      1      2      3      4
4     -1      1      3      5
5 >> A3 = [1 3
6          -4 7]
7 A3 =
8      1      3
9     -4      7

```

From that point of view, a row vector is a $1 \times k$ matrix and a column vector is an $n \times 1$ matrix. Transposing a vector changes it from a row to a column or the other way around. This idea can be extended to a matrix, where the transpose operation interchanges rows with the corresponding columns, as in the example:

```

1 >> A2
2 A2 =
3      1      2      3      4
4     -1      1      3      5
5 >> A2'                                % transpose of A2
6 ans =
7      1     -1
8      2      1
9      3      3
10     4      5
11 >> size(A2)                          % returns the size (dimensions) of A2: 2 rows,
12                                         % 4 columns
13 ans =
14      2      4
15 >> size(A2')
16 ans =
17      4      2

```

13.2.1 Special matrices

There is a number of built-in matrices of size specified by the user (see Table 13.2). A few examples are given below:

```

1 >> E = []                             % an empty matrix of 0-by-0 elements!
2 E =
3      []
4 >> size(E)
5 ans =
6      0      0
7 >> I = eye(3);                         % the 3-by-3 identity matrix
8 I =
9      1      0      0
10     0      1      0
11     0      0      1
12 >> r = [1 3 -2]; R = diag(r)          % create a diagonal matrix with r on the diagonal
13 R =
14      1      0      0
15     0      3      0

```

```

16      0      0      -2
17 >> A = [1 2 3; 4 5 6; 7 8 9];
18 >> diag(A) % extracts the diagonal entries of A
19 ans =
20      1
21      5
22      9
23 >> B = ones(3,2)
24 B =
25      1      1
26      1      1
27      1      1
28 >> C = zeros (size(B')) % a matrix of all zeros of the size given by B'
29 C =
30      0      0      0
31      0      0      0
32 >> D = rand(2,3) % a matrix of random numbers; you will get a
33 % different one!
34 D =
35      0.0227      0.9101      0.9222
36      0.0299      0.0640      0.3309
37 >> v = linspace(1,2,4) % a vector is also an example of a matrix
38 v =
39      1.0000      1.3333      1.6667      2.0000

```

13.2.2 Building matrices and extracting parts of matrices

It is often needed to build a larger matrix from a number of smaller ones:

```

1 >> x = [4; -1], y = [-1 3]
2 x =
3      4
4     -1
5 y =
6     -1      3
7 >> X = [x y'] % X consists of the columns x and y'
8 X =
9      4     -1
10     -1      3
11 >> T = [ -1 3 4; 4 5 6]; t = 1:3;
12 >> T = [T; t] % add to T a new row, namely the row vector t
13 T =
14     -1      3      4
15      4      5      6
16      1      2      3
17 >> G = [1 5; 4 5; 0 2]; % G is a matrix of the 3-by-2 size; check size(G)
18 >> T2 = [T G] % join two matrices
19 T2 =
20     -1      3      4      1      5
21      4      5      6      4      5
22      1      2      3      0      2
23 >> T3 = [T; G ones(3,1)] % G is 3-by-2, T is 3-by-3
24 T3 =
25     -1      3      4
26      4      5      6
27      1      2      3
28      1      5      1
29      4      5      1

```

```

30      0      2      1
31 >> T3 = [T; G']; % this is also possible
32 T3 =
33     -1      3      4
34      4      5      6
35      1      2      3
36      1      4      0
37      5      5      2
38 >> [G' diag(5:6); ones(3,2) T] % you can join many matrices
39 ans =
40      1      4      0      5      0
41      5      5      2      0      6
42      1      1     -1      3      4
43      1      1      4      5      6
44      1      1      1      2      3

```

A part of a matrix can be extracted from a matrix in a similar way as it is done for vectors. Each element in a matrix is indexed by a row and a column to which it belongs. Mathematically, the element from the i -th row and the j -th column of the matrix A is denoted as A_{ij} ; MATLAB provides the $A(i,j)$ notation.

```

1 >> A = [1:3; 4:6; 7:9]
2 A =
3      1      2      3
4      4      5      6
5      7      8      9
6 >> A(1,2), A(2,3), A(3,1)
7 ans =
8      2
9 ans =
10     6
11 ans =
12     7
13 >> A(4,3) % this is not possible: A is a 3-by-3 matrix!
14 ??? Index exceeds matrix dimensions.
15 >> A(2,3) = A(2,3) + 2*A(1,1) % change the value of A(2,3)
16 A =
17      1      2      3
18      4      5      8
19      7      8      9

```

It is easy to automatically extend the size of a matrix. For the matrix A above it can be done e.g. as follows:

```

1 >> A(5,2) = 5 % assign 5 to the position (5,2); the uninitialized
2 A = % elements of A become zeros
3      1      2      3
4      4      5      8
5      7      8      9
6      0      0      0
7      0      5      0

```

If needed, the other zero elements of the matrix A can also be defined, by e.g.:

```

1 >> A(4,:) = [2, 1, 2]; % assign vector [2, 1, 2] to the 4th row of A
2 >> A(5,[1,3]) = [4, 4]; % assign: A(5,1) = 4 and A(5,3) = 4

```

```

3 >> A                                % what does the matrix A look like now?
4 A =
5     1     2     3
6     4     5     8
7     7     8     9
8     2     1     2
9     4     5     4

```

Different parts of the matrix A can now be extracted:

```

1 >> A(3,:)                            % extract the 3rd row of A
2 ans =
3     7     8     9
4 >> A(:,2)                            % extract the 2nd column of A
5 ans =
6     2
7     5
8     8
9     1
10    5
11 >> A(1:2,:)                          % extract the 1st and 2nd row of A
12 ans =
13     1     2     3
14     4     5     8
15 >> A([2,5],1:2)                      % extract a part of A
16 ans =
17     4     5
18     4     5

```

As you have seen in the examples above, it is possible to manipulate (groups of) matrix-elements. The commands are shortly explained in Table 13.1.

The concept of an empty matrix `[]` is also very useful in MATLAB. For instance, a few columns or rows can be removed from a matrix by assigning an empty matrix to it. Try for example:

```

1 >> C = [1 2 3 4; 5 6 7 8; 1 1 1 1];
2 >> D = C; D(:,2) = []                % now a copy of C is in D; remove the 2nd column
3                                     % of D
4 >> C ([1,3],:) = []                  % remove the rows 1 and 3 from C

```

13.2.3 Operations on matrices

Table 13.2 shows some frequently used matrix operations and functions. The important ones are dot operations on matrices, matrix-vector products and matrix-matrix products. In the class of the dot operations, there are dot product, dot division and dot power. Those operations work as for vectors: they address matrices in the element-by-element way, therefore they can be performed on matrices of the same sizes. They also allow for scalar-matrix operations. For the dot product or division, the corresponding elements are either multiplied or divided. A few examples of basic operations are given below:

```

1 >> B = [1 -1 3; 4 0 7]
2 B =
3     1    -1     3

```

```

4      4      0      7
5 >> B2 = [1 2; 5 1; 5 6];
6 >> B = B + B2' % add two matrices; why is B2' needed instead of B2?
7 B =
8      2      4      8
9      6      1     13
10 >> B-2 % subtract 2 from all elements of B
11 ans =
12      0      2      6
13      4     -1     11
14 >> ans = B./4 % divide all elements of the matrix B by 4
15 ans =
16      0.5000      1.0000      2.0000
17      1.5000      0.2500      3.2500
18 >> 4/B % this is not possible
19 ??? Error using ==> /
20 Matrix dimensions must agree.
21
22 >> 4./B % this is possible;
23 % equivalent to: 4.*ones(size(B))./ B
24 ans =
25      2.0000      1.0000      0.5000
26      0.6667      4.0000      0.3077
27 >> C = [1 -1 4; 7 0 -1];
28 >> B .* C % multiply element-by-element
29 ans =
30      2     -4     32
31     42      0    -13
32 >> ans.^3 - 2 % do for all elements: raise to the power 3 and
33 % subtract 2
34 ans =
35      6     -66    32766
36    74086      -2    -2199
37 >> ans ./ B.^2 % element-by-element division
38 ans =
39      0.7500     -1.0312    63.9961
40    342.9907     -2.0000    -1.0009
41 >> r = [1 3 -2]; r * B2 % this is a legal operation:
42 ans = % r is a 1-by-3 matrix and B2 is a 3-by-2 matrix
43      6     -7 % note that B2 * r is an illegal operation

```

Concerning the matrix-vector and matrix-matrix products, two things should be reminded from linear algebra. First, an $n \times k$ matrix A (having n rows and k columns) can be multiplied by a $k \times 1$ (column) vector x , resulting in a column $n \times 1$ vector y , i.e.: $Ax = y$ such that $y_i = \sum_{p=1}^k A_{ip} x_p$. Multiplying a $1 \times n$ (row) vector x by a matrix A , results in a $1 \times k$ (row) vector y . Secondly, an $n \times k$ matrix A can be multiply by a matrix B , only if B has k rows, i.e. B is $k \times m$ (m is arbitrary). As a result, you get $n \times m$ matrix C , such that $AB = C$, where $C_{ij} = \sum_{p=1}^k A_{ip} B_{pj}$.

```

1 >> b = [1 3 -2];
2 >> B = [1 -1 3; 4 0 7]
3 B =
4      1     -1      3
5      4      0      7
6 >> b * B % not possible: b is 1-by-3 and B is 2-by-3
7 ??? Error using ==> *
8 Inner matrix dimensions must agree.
9
10 >> b * B' % this is possible: a row vector multiplied by a

```

```

11                                     % matrix
12 ans =
13     -8    -10
14 >> B' * ones(2,1)
15 ans =
16     5
17    -1
18    10
19 >> C = [3 1; 1 -3];
20 >> C * B
21 ans =
22     7    -3    16
23    -11    -1   -18
24 >> C.^3                                     % this is an element-by-element power
25 ans =
26    27     1
27     1   -27
28 >> C^3                                     % this is equivalent to C*C*C
29 ans =
30    30    10
31    10   -30
32 >> ones(3,4)./4 * diag(1:4)
33 ans =
34    0.2500    0.5000    0.7500    1.0000
35    0.2500    0.5000    0.7500    1.0000
36    0.2500    0.5000    0.7500    1.0000

```

13.3 Exercises

Exercise 13.2.

Create a vector containing 5 elements such that its components are equally spaced when you take the logarithm of it ■

Exercise 13.3.

Perform the following exercises:

- Create a vector \mathbf{x} with the elements:
 - 1, $1/2$, $1/3$, $1/4$, $1/5$
 - 0, $1/2$, $2/3$, $3/4$, $4/5$

To do this, divide a vector \mathbf{y} by a vector \mathbf{z} .

- Create a vector \mathbf{x} with the elements: $x_n = \frac{(-1)^n}{2n-1}$ for $n = 1, 2, 3, \dots$. Find the sum of the 100-element vector.
- Given a vector \mathbf{t} , write down the MATLAB expressions that will compute:
 - $\ln(2 + t + t^2)$
 - $\cos(t)^2 - \sin(t)^2$
 - $e^t(1 + \cos(3t))$
 - $\tan^{-1}(t)$

Test them for $\mathbf{t} = 1 : 0.2 : 2$.

■

Exercise 13.4.

Use the knowledge on computing the inner product to find:

1. the Euclidean length of the vector $\mathbf{x} = [2 \ 1 \ 3 \ 7 \ 9 \ 4 \ 6]$, which is defined as $\|\mathbf{x}\| = \sqrt{(\Sigma \mathbf{x}_i^2)}$.
2. the angle between two column vectors, which is defined as $\cos \alpha = \frac{\mathbf{x}^T \mathbf{y}}{\|\mathbf{x}\| \|\mathbf{y}\|}$. Note that you can also use the MATLAB command `norm(v)`, which gives you the Euclidean length of a vector v . The MATLAB command `acosd(x)` can be used to find the inverse cosine of x expressed in degrees. Compute the angle between two vectors:
 - $\mathbf{x} = [3 \ 2 \ 1]$ and $\mathbf{y} = [1 \ 2 \ 3]$
 - $\mathbf{x} = 1 : 5$ and $\mathbf{y} = 6 : 10$

■

Exercise 13.5.

Clear all variables (use the command `clear`). Define the matrix $\mathbf{A} = [1:4; 5:8; 1 \ 1 \ 1 \ 1]$. Predict and check the result of the following operations:

- $\mathbf{x} = \mathbf{A}(:, 3)$
- $\mathbf{B} = \mathbf{A}(1 : 3, 2 : 2)$
- $\mathbf{A}(1, 1) = 9 + \mathbf{A}(2, 3)$
- $\mathbf{A}(2 : 3, 1 : 3) = [0 \ 0 \ 0; 0 \ 0 \ 0]$
- $\mathbf{A}(2 : 3, 1 : 2) = [1 \ 1; 3 \ 3]$
- $\mathbf{y} = \mathbf{A}(3 : 3, 1 : 4)$
- $\mathbf{A} = [\mathbf{A}; 2 \ 1 \ 7 \ 7; 7 \ 7 \ 4 \ 5]$
- $\mathbf{C} = \mathbf{A}([1, 3], 2)$
- $\mathbf{D} = \mathbf{A}([2, 3, 5], [1, 3, 4])$
- $\mathbf{D}(2, :) = []$

■

Exercise 13.6.

Define the matrices $\mathbf{T} = [\ 3 \ 4; \ 1 \ 8; \ -4 \ 3]$ and $\mathbf{A} = [\text{diag}(-1:2:3) \ \mathbf{T}; \ -4 \ 4 \ 1 \ 2 \ 1]$. Perform the following operations on the matrix \mathbf{A} :

- extract a vector consisting of the 2nd and 4th elements of the 3rd row
- find the minimum of the 3rd column
- find the maximum of the 2nd row
- compute the sum of the 2nd column
- compute the mean of the 1st and 4th rows
- extract the submatrix consisting of the 1st and 3rd rows and all columns
- extract the submatrix consisting of the 1st and 2nd rows and the 3rd, 4th and 5th columns
- compute the total sum of the 1st and 2nd rows
- add 3 to all elements of the 2nd and 3rd columns

■

Exercise 13.7.

Let $\mathbf{A} = [2 \ 4 \ 1; \ 6 \ 7 \ 2; \ 3 \ 5 \ 9]$. Provide the commands which:

- assign the first row of \mathbf{A} to a vector \mathbf{x} ;
- assign the last 2 rows of \mathbf{A} to a vector \mathbf{y} ;
- add up the columns of \mathbf{A} ;

- add up the rows of A;

■

Exercise 13.8.

Let $A = \begin{bmatrix} 2 & 7 & 9 & 7; & 3 & 1 & 5 & 6; & 8 & 1 & 2 & 5 \end{bmatrix}$. Explain the results or perform the following commands:

- | | | |
|-----------------------------|-------------------------|---------------------------------------------------------------|
| • A' | • $\text{sum}(A)$ | • $\max(\min(A))$ |
| • $A(1,:)'$ | • $\text{sum}(A')$ | • $[[A; \text{sum}(A)] [\text{sum}(A, 2); \text{sum}(A(:))]]$ |
| • $A(:, [14])$ | • $\text{mean}(A)$ | • assign the even-numbered columns of A to an array B |
| • $A([23], [31])$ | • $\text{mean}(A')$ | • assign the odd-numbered rows to an array C |
| • $\text{reshape}(A, 2, 6)$ | • $\text{sum}(A, 2)$ | • convert A into a 4-by-3 array |
| • $A(:)$ | • $\text{mean}(A, 2)$ | • compute the reciprocal of each element of A |
| • $\text{flipud}(A)$ | • $\min(A)$ | • compute the square-root of each element of A |
| • $\text{fliplr}(A)$ | • $\max(A')$ | • remove the second column of A |
| • $[A \ A(\text{end}, :)]$ | • $\min(A(:, 4))$ | • add a row of all 1's at the beginning and at the end |
| • $[A; A(1:2, :)]$ | • $[\min(A)' \max(A)']$ | • swap the 2nd row and the last row |

■

Exercise 13.9.

Given the vectors $x = [1 \ 3 \ 7]$, $y = [2 \ 4 \ 2]$ and the matrices $A = \begin{bmatrix} 3 & 1 & 6; & 5 & 2 & 7 \end{bmatrix}$ and $B = \begin{bmatrix} 1 & 4; & 7 & 8; & 2 & 2 \end{bmatrix}$, determine which of the following statements can be correctly executed (and if not, try to understand why) and provide the result:

- | | | | |
|-------------------|-------------|-------------|---------------------------|
| • $x + y$ | • $[x; y']$ | • $B * A$ | • $B ./ x'$ |
| • $x + A$ | • $[x; y]$ | • $A .* B$ | • $B ./ [x' \ x']$ |
| • $x' + y$ | • $A - 3$ | • $A' .* B$ | • $2/A$ |
| • $A - [x' \ y']$ | • $A + B$ | • $2 * B$ | • $\text{ones}(1, 3) * A$ |
| • $[x; y] + A$ | • $B' + A$ | • $2. * B$ | • $\text{ones}(1, 3) * B$ |

■

Exercise 13.10.

Perform all operations from Table 13.2, using some matrices A and B, vector v and scalars k, a, b, n, and m. ■

Exercise 13.11.

Let A be a square 6-by-6 matrix.

1. Create a matrix B, whose elements are the same as those of A except the entries on the main diagonal. The diagonal of B should consist of 1s.
2. Create a tridiagonal matrix T, whose three diagonal are taken from the matrix A. *Hint:* you may use the commands `triu` and `tril`.

■

Exercise 13.12.

Let A be a random 5×5 matrix and let b be a random 5×1 vector. Given that $Ax = b$, try to find x (look at Table 13.2). Explain what is the difference between the operators \backslash , $/$ and the command `inv`. Having found x , check whether $Ax - b$ is close to a zero vector. ■

Exercise 13.13.

Let $A = \text{ones}(6) + \text{eye}(6)$. Normalize the columns of the matrix A so that all columns of the resulting matrix, say B , have the Euclidean norm (length) equal to 1. Next, find the angles between consecutive columns of the matrix B . ■

Exercise 13.14.

Find two 2×2 matrices A and B for which $A.*B \neq A*B$ holds. Make use of the following operations: $/$, \backslash , or the command `inv`. ■

Chapter 14

Scripts

14.1 Script m-files

MATLAB commands can be entered at the MATLAB prompt. When a problem is more complicated this becomes inefficient. A solution is to use script m-files.

M-files are useful when the number of commands increases or when you want to change values of some variables and re-evaluate them quickly. Formally, a script is an external file that contains a sequence of MATLAB commands (statements). However, it is not a function, since there are no input/output parameters and the script variables remain in the workspace. So, when you run a script, the commands in it are executed as if they have been entered through the keyboard.

The use of m-files is illustrated in the following paragraph.

How to create and run a script

- Open the MATLAB editor (go to the File menu-bar, choose the New option followed by m-file; as a result, the MATLAB Editor Window will appear)
- Enter the lines listed below

```
1      x = 0:0.2:6;  
2      y = sin(x);  
3      plot(x,y);  
4      title('Plot of y = sin(x)');
```

- Save the files as sinplot.m
- Run the script by typing in the MATLAB prompt:

```
1 sinplot
```

Alternatively to run the script you can press F5 or you can click on the run icon.

You might have problems with this because your path is not set correctly. The path is a collection of directories (folders) where MATLAB goes to look for programs to be run. If you saved that file first in, say, Myfolder, you should tell this to MATLAB by typing:

```
1 >> cd Myfolder
```

or opening the path browser utility from your MATLAB window.

Note that the sinplot script affects the workspace. Check:

```
1 clear % all variables are removed from the workspace
2 who % no variables present
3 sinplot
4 who
5 Your variables are: x y
```

These generic names, x and y, may easily be used in further computations and this can cause side effects. Side effects occur in general when a set of commands change variables other than the input arguments. Since scripts create and change variables in the workspace (without warning), a bug, hard to track down, may easily appear. So, it is important to remember that the commands within a script have access to all variables in the workspace and all variables created in this script become a part of the workspace. Therefore, it is better to use function m-files to solve a specific problem.

14.2 Exercises

Exercise 14.1.

Write a script that for the given a , b and c returns the roots of the cubic equation: $ax^2+bx+c=0$. (Hint: the solution to the equation is given by $x = \frac{-b \pm \sqrt{b^2-4ac}}{2a}$) ■

Chapter 15

Visualization

MATLAB includes very powerful features for easy figure creation, plotting and image display. MATLAB can be used to visualize the results of an experiment. We will start by introducing the most basic features here.

15.1 2D plots

The most important commands for basic 2-dimensional plotting are illustrated in the following example:

Example 15.1:

Plotting the sinus function

```
1 x=0:0.1:10; % x ranges from 0 to 10 in steps of 0.1
2 y=sin(x);   % y contains the sin of each value of x
3 figure      % create an empty figure window
4 plot(y)     % plots all of the y values (the values on the x-axis are
5             % indexes from 1 to 101)
6 figure      % create an empty figure window
7 plot(x,y)   % plots y against x (replacing the old plot on the same
8             % figure, now showing the units of x on the x-axis)
9 close       % closes the figure window
```


Note that `x` and `y` have to be both either row or column vectors of the same length (i.e. having the same number of elements).

To plot a graph of a function, it is important to sample the function sufficiently well. Compare the following examples:

```

1      % coarse sampling
2      n = 5;
3      x = 0:1/n:3;
4      y = sin(5*x);
5      plot(x,y)


```



```

1      % good sampling
2      n = 25;
3      x = 0:1/n:3;
4      y = sin(5*x);
5      plot(x,y)

```



The `plot` command has many additional possible arguments, as can be seen by typing `help plot`. There are also many other plotting commands for 2-dimensional plotting (`help graph2d`) and also for 3-D plotting (`help graph3d`). You can close a figure window by clicking on the upper right X box, using the `close` command on the figure's file menu, or by typing `close` at the command prompt. Use `close all` to close all figures and start a new task (or use `close 1` to close Figure no.1 etc).

If you do not create an extra figure using `figure`, the previous function is removed as soon as the next is displayed.

The color and point marker can be changed on a plot by adding a third parameter (in single quotes) to the plot command. For example, to plot the a function as a red, dotted line, the m-file should be changed to:

```

1      x = 0:0.1:100;
2      y = 3*x;
3      plot(x,y, 'r: ')

```

The third input consists of one to three characters which specify a color and/or a point marker type. The list of colors and styles is reported in Table 15.1.

The commands `loglog`, `semilogx` and `semilogy` are similar to `plot`, except that they use either one or two logarithmic axes.

15.2 Several functions in one figure

You can plot more than one function on the same figure. To plot a sine wave and cosine wave on the same set of axes, using a different color and point marker for each, the following m-file

Symbol	Color	Symbol	Line style
r	red	., o	point, circle
g	green	*	star
b	blue	x, +	x-mark, plus
y	yellow	-	solid line
m	magenta	--	dash line
c	cyan	:	dot line
k	black	-.	dash-dot line

Table 15.1: Plot colors and styles.

could be used:

```

1      x = 1:0.1:2*pi;
2      y = sin(x);
3      z = cos(x);
4      plot(x,y,'r', x,z,'gx')
```

By adding more sets of parameters to plot, you can plot as many different functions on the same figure as you want. When plotting many things on the same graph it is useful to differentiate the different functions based on color and point marker. This same effect can also be achieved using the `hold on` and `hold off` commands. The same plot shown above could be generated using the following m-file:

```

1      x = linspace(0,2*pi,50);
2      y = sin(x);
3      z = cos(x);
4      hold on
5      plot(x,y,'r')
6      plot(x,z,'gx')
7      hold off
```

Always remember that if you use the `hold on` command, all plots from then on will be generated on one set of axes, without erasing the previous plot, until the `hold off` command is issued.

It is also possible to produce a few subplots in one figure window. With the command `subplot`, the window can be horizontally and vertically divided into $p \times r$ subfigures, which are counted from 1 to pr , row-wise, starting from the top left. The commands: `plot`, `title`, `grid` etc work only in the current subfigure. So, if you want to change something in other subfigure, use the command `subplot` to switch there.

```

1      x = 1:.1:4;
2      y1 = sin(3*x);
3      y2 = cos(5*x);
4      y3 = sin(3*x).*cos(5*x);
5      figure
6      subplot(1,3,1);
7      plot(x,y1,'m');
8      title('sin(3*x)')
9      subplot(1,3,2);
10     plot(x,y2,'g');
```

```
11     title('cos(5*x)')
12     subplot(1,3,3);
13     plot(x,y3,'k');
14     title('sin(3*x) * cos(5*x)')
```

15.3 Adding text

Another thing that may be important for your plots is labeling. You can give your plot a title (with the `title` command), x-axis label (with the `xlabel` command), y-axis label (with the `ylabel` command), and put text on the actual plot. All of the above commands are issued after the actual plot command has been issued.

Furthermore, text can be put on the plot itself in one of two ways: the `text` command and the `gtext` command. The first command involves knowing the coordinates of where you want the text string. The command is `text(xcor,ycor,'textstring')`. To use the other command, you do not need to know the exact coordinates. The command is `gtext('textstring')`, and then you just move the cross-hair to the desired location with the mouse, and click on the position you want the text placed.

To add a title, grid and to label the axes, one uses:

```
1     x=0:0.1:3;
2     y=sin(x);
3     z=cos(x);
4     hold on;
5     plot(x,y);
6     plot(x,z);
7     title('Sine and Cosine')
8     xlabel('x')
9     ylabel('sin(x) and cos(x)')
10    gtext('unnecessary labeling')
11    legend('sin','cos')
```

The text "unnecessary labeling" was placed right above the position, I clicked on.

Table 15.2 shows a few possibilities of the `plot` command, `help plot` shows them all.

15.4 Editing plots

MATLAB formats a graph to provide readability, setting the scale of axes, including tick marks on the axes, and using color and line style to distinguish the plots in the graph. However, if you are creating presentation graphics, you may want to change this default formatting or add descriptive labels, titles, legends and other annotations to help explain your data. MATLAB supports two ways to edit the plots you create:

- Using the Property Editor to select and edit objects interactively.

Command	Result
<code>grid on/off</code>	adds a grid to the plot at the tick marks or removes it
<code>box off/on</code>	removes the axes box or shows it
<code>axis([xmin xmax ymin ymax])</code>	sets the minimum and maximum values of the axes
<code>xlabel('text')</code>	plots the label text on the x-axis
<code>ylabel('text')</code>	plots the label text on the y-axis
<code>zlabel('text')</code>	plots the label text on the z-axis
<code>title('text')</code>	plots a title above the graph
<code>text(x,y,'text')</code>	adds text at the point (x,y)
<code>gtext('text')</code>	adds text at a manually (with a mouse) indicated point
<code>legend('fun1','fun2')</code>	plots a legend box (move it with your mouse) to name your functions
<code>legend off</code>	deletes the legend box
<code>clf</code>	clear the current figure
<code>figure(n)</code>	creates a figure number <i>n</i>
<code>subplot</code>	creates a subplot in the current figure

Table 15.2: Useful commands to make plots.

- Using MATLAB functions at the command-line or in an M-file (look at "Line Properties" in the help for the list of properties).

```

1    figure
2    hold on
3    x = 1:.1:4;
4    h1=plot(x,cos(x))
5    h=plot(x,sin(x))
6    set(h1,'LineWidth',2,'LineStyle','.', 'Color','g')
7    set(h,'LineWidth',2,'LineStyle','^', 'Color','r')
8    hold off

```

In this above example, **h** and **h1** are handles (basically pointers to objects, see c++ part). The function `get(h)` returns all properties of the handle **h**, this can be very useful for touching up plots.

15.5 Changing the axis

An important way to customize your plots to meet your needs is with the `axis` command. The `axis` command changes the axis of the plot shown, so only the part of the axis that is desirable is displayed. The `axis` command is used by entering the following command right after the plot command (or any command that has a plot as an output):

```

1    axis([xmin, xmax, ymin, ymax])

```


15.6 Exporting graph

To copy the plot into Word you can select directly on the figure

```
1 Edit → Copy Figure
```

You can also print to a file if you specify the file name. If you do not provide an extension, `print` adds one. Since there are many parameters they will not be explained here (check `help print` to learn more). Instead, try to understand the examples:

```
1 % print the current Figure to the current printer in color
2 print -dwinc
3 % print Figure no.1 to the file myfile.eps in black
4 print -fl -deps myfile.eps
5 % print Figure no.1 to the file myfilec.eps in color
6 print -fl -depsc myfilec.eps
7 % print the current Figure to the file myfile1.tiff
8 print -dtiff myfile1.tiff
9 % print the current Figure to the file myfile1.ps in color
10 print -dpsc myfile1c.ps
11 % print Figure no.2 to the file myfile2.jpg
12 print -f2 -djpeg myfile2
```

15.7 Plotting surfaces

MATLAB provides a number of commands to plot 3D data. A surface is defined by a function $f(x, y)$, where for each pair of (x, y) , the height z is computed as $z = f(x, y)$. To plot a surface, a rectangular domain of the (x, y) -plane should be sampled. The mesh (or grid) is constructed by the use of the command `meshgrid` as follows:

```
1 [X, Y] = meshgrid (-1:.5:1, 0:.5:2)
2 X =
3 -1.0000 -0.5000 0 0.5000 1.0000
4 -1.0000 -0.5000 0 0.5000 1.0000
5 -1.0000 -0.5000 0 0.5000 1.0000
6 -1.0000 -0.5000 0 0.5000 1.0000
7 -1.0000 -0.5000 0 0.5000 1.0000
8 Y =
9 0 0 0 0 0
10 0.5000 0.5000 0.5000 0.5000 0.5000
11 1.0000 1.0000 1.0000 1.0000 1.0000
12 1.5000 1.5000 1.5000 1.5000 1.5000
13 2.0000 2.0000 2.0000 2.0000 2.0000
```

The domain $[-1, 1] \times [0, 2]$ is now sampled with 0.5 in both directions and it is described by points $[X(i, j), Y(i, j)]$. To plot a smooth surface, the chosen domain should be sampled in a dense way. To plot a surface, the command `mesh` or `surf` can be used:

```

1 [X,Y] = meshgrid(-1:.05:1, 0:.05:2);
2 Z = sin(5*X) .* cos(2*Y);
3 mesh(X,Y,Z);
4 title ('Function z = sin(5x) * cos(2y)')

```

Use colormap to define different colors for plotting.

15.7.1 Contour plots

A filled (2D) contour plot displays isolines calculated from matrix z and fills the areas between the isolines using constant colors. The color of the filled areas depends on the current figure's colormap.

Example 15.2:

To view a contour plot of the function $z = x \exp(-x^2 - y^2)$ over the range $-2 \leq x \leq 2$ and $-2 \leq y \leq 3$

```

1 [X,Y] = meshgrid(-2:.05:2,-2:.05:3);
2 Z = X.*exp(-X.^2-Y.^2);
3 figure
4 [C,h]=contourf(X,Y,Z,25);
5 colormap hsv
6 figure
7 [C,h] = contour(X,Y,Z,15);
8 close

```





Figure 15.1: Output of the contourf and contour example

Example 15.3:

To locate e.g. the minimum value of the function the function $z = x \exp(-x^2 - y^2)$ on the grid, you can proceed as follows:

```

1 [X,Y] = meshgrid(-2:.05:2,-2:.05:3);
2 Z = X.*exp(-X.^2-Y.^2);
3 [mm,I] = min(Z); % a row vector of the min. elements from each column

```

```

4           % I is a vector of corresponding
5   [Zmin, j] = min (mm); % Zmin is the minimum value, j is the index
6           % Zmin is the value of  Z(I(j),j)
7   xpos = X(I(j),j); %
8   ypos = Y(I(j),j); % position of the minimum value
9   contour (X,Y,Z,25);
10  xlabel('x-axis'); ylabel('y-axis');
11  hold on
12  plot(xpos(1),ypos,'*');
13  text(xpos(1)+0.1,ypos,'Minimum');
14  hold off

```

15.8 3D line plots

The command `plot3` to plot lines in 3D is equivalent to the command `plot` in 2D. The format is the same as for `plot`, it is, however, extended by an extra coordinate. An example is plotting the curve r defined parametrically as $r(t) = [t \sin(t), t \cos(t), t]$ over the interval $[-10\pi, 10\pi]$.

```

1   t = linspace(-10*pi,10*pi,200);
2   plot3(t.*sin(t), t.*cos(t), t, 'md-'); % plot the curve in magenta
3   title('Curve r(t) = [t sin(t), t cos(t), t]');
4   xlabel('x-axis');
5   ylabel('y-axis');
6   zlabel('z-axis');
7   grid

```

15.9 Animations

A sequence of graphs can be put in motion in MATLAB (the version should be at least 5.0), i.e. you can make a movie using MATLAB graphics tools. The MATLAB command `addframe(avifile,frame)` appends the data in `frame` to the AVI file identified by the variable `avifile`, which was created by a previous call to the file `output.avi`.

Example 15.4:

Create a movie of an evolving sinus curve in the interval 4π .

```

1   clear all
2   fig=figure;
3   mov = avifile('output.avi')
4   x = 0:0.1:4*pi;
5   xx(1) = 0;
6   yy(1) = 0;
7   for k=1:length(x)
8       xx(k) = x(k);
9       yy(k) = sin(x(k));
10      h=plot(xx,yy);
11      axis([0 4*pi -1.5 1.5])
12      set(h,'color','m','linewidth',2,'marker','o')
13      F=getframe(fig);
14      mov = addframe(mov,F);

```

```

15     end
16     mov = close(mov);

```

Here, a `for`-loop construction has been used to create the movie frames. You will learn more on loops in the next classes, see section 16.3.

Via the command `getframe` each frame is stored in the column of the matrix `mov`. Note that to create a movie requires quite some memory.

15.10 Exercises

Exercise 15.1.

Make a plot connecting the coordinates: (2, 6), (2.5, 18), (5, 17.5), (4.2, 12.5) and (2,12) by a line. ■

Exercise 15.2.

Plot the function $y = \sin(x) + x - x \cos(x)$ in two separate figures for the intervals: $0 < x < 30$ and $-100 < x < 100$. Add a title and axes description. ■

Exercise 15.3.

Plot a circle with the radius $r = 2$, knowing that the parametric equation of a circle is $[x(t), y(t)] = [r \cos(t), r \sin(t)]$ for $t = [0, 2\pi]$. ■

Exercise 15.4.

Plot an ellipse with semiaxes $a = 4$ and $b = 2$. ■

Exercise 15.5.

Plot the functions $f(x) = x$, $g(x) = x^3$, $h(x) = e^x$ and $z(x) = e^{x^2}$ over the interval $[0, 4]$ on the normal scale and on the log-log scale. Use an appropriate sampling to get smooth curves. Describe your plots by using the functions: `xlabel`, `ylabel`, `title` and `legend`. ■

Exercise 15.6.

Make a plot of the functions: $f(x) = \sin(1/x)$ and $f(x) = \cos(1/x)$ over the interval $[0.01, 0.1]$. How do you create x so that the plots look sufficiently smooth? ■

Exercise 15.7.

Produce a nice graph which demonstrates as clearly as possible the behavior of the function $f(x, y) = \frac{xy^2}{x^2 + y^4}$ near the point (0,0). Note that the sampling around this points should be dense enough. ■

Exercise 15.8.

Plot a sphere, which is parametrically defined as $[x(t, s), y(t, s), z(t, s)] = [\cos(t) * \cos(s), \cos(t) * \sin(s), \sin(t)]$ for $t, s = [0, 2\pi]$ (use `surf`). Make first equal axes, then remove them. Use shading `interp` to remove black lines (use `shading faceted` to restore the original picture). ■

Exercise 15.9.

Plot the parametric function of r and θ : $[x(r, \theta), y(r, \theta), z(r, \theta)] = [r \cos(\theta), r \sin(\theta), \sin(6 \cos(r) - n\theta)]$ for $\theta = [0, 2\pi]$ and $r = [0, 4]$. Choose n to be constant. Observe, how the graph changes depending on different n . ■

Exercise 15.10.

Plot the surface $f(x, y) = xy e^{-x^2-y^2}$ over the domain $[-2, 2] \times [-2, 2]$. Find the values and the locations of the minima and maxima of this function. ■

Exercise 15.11.

Make a 3D smooth plot of the curve defined parametrically as: $[x(t), y(t), z(t)] = [\sin(t), \cos(t), \sin^2(t)]$ for $t = [0, 2\pi]$. Plot the curve in green, with the points marked by circles. Add a title, description of axes and the grid. You can rotate the image by clicking `Tools` at the Figure window and choosing the `Rotate 3D` option or by typing `rotate3D` at the prompt. Then by clicking at the image and dragging your mouse you can rotate the axes. Experiment with this option. ■

Exercise 15.12.

Write a script that makes a movie consisting of 5 frames of the surface $f(x, y) = \sin(nx) \sin(ky)$ over the domain $[0, 2\pi] \times [0, 2\pi]$ and $n = 1 : 5$. Add a title, description of axes and shading. ■

Exercise 15.13.

Plot an equilateral triangle with two vertices $[a \ a]$ and $[b \ a]$. Find the third vertex. Use `fill` to paint the triangle. ■

Chapter 16

Control flow

A control flow structure is a block of commands that allows conditional code execution and making loops.

16.1 Logical and relational operators

To use control flow commands, it is necessary to perform operations that result in logical values: TRUE or FALSE. In MATLAB the result of a logical operation is 1 if it is true and 0 if it is false. Table 16.1 shows the relational and logical operations. Another way to get to know more about them is to type `help relop`. The relational operators `<`, `<=`, `>`, `>=`, `==` and `~=` can be used to compare two arrays of the same size or an array to a scalar. The logical operators `&`, `|` and `~` allow for the logical combination or negation of relational operators. In addition, three functions are also available: `xor`, `any` and `all` (use `help` to find out more).

Important: The logical `&` and `|` have the equal precedence in MATLAB, which means that those operators associate from left to right. A common situation is:

```
1 >> b = 10;
2 >> 1 | b > 0 & 0
3 ans =
4      0
5 >> (1 | b > 0) & 0           % this indicates the same as above
6 ans =
7      0
8 >> 1 | (b > 0 & 0)
9 ans =
10     1
```

Command	Result
<code>a = (b > c)</code>	a is 1 if b is larger than c. Similar are: <code><</code> , <code>>=</code> and <code><=</code>
<code>a = (b == c)</code>	a is 1 if b is equal to c
<code>a = (b ~= c)</code>	a is 1 if b is not equal c
<code>a = ~b</code>	logical complement: a is 1 if b is 0
<code>a = (b & c)</code>	logical AND: a is 1 if b = TRUE AND c = TRUE
<code>a = (b c)</code>	logical OR: a is 1 if b = TRUE OR c = TRUE

Table 16.1: Relational and logical operations.

This shows that you should *always* use brackets to indicate in which way the operators should be evaluated.

The introduction of the logical data type has forced some changes in the use of non-logical 0-1 vectors as indices for subscripting. You can see the differences by executing the following commands that attempt to extract the elements of `y` that correspond to either the odd or even elements of `x`, assuming that `x` and `y` are two vectors of the same length.

- `y(rem(x,2))` vs. `y(logical(rem(x,2)))` % odd elements
- `y(~rem(x,2))` vs. `y(~logical(rem(x,2)))` % even elements

These examples show that a *numerical* 0 can not be used as an array index, i.e. the first expression will result in a MATLAB error if any of the remainders computed with the `rem` function equals zero. Converting these zeros into *logical* 0's solves the problem. Note that both expressions to select the even elements do work as the logical complement `~` implicitly converts the numerical results of `rem` into logical values.

Exercise 16.1.

Exercise with logical and relational operators:

1. Predict and check the result of each of the operations of Table 16.1 for `b = 0` and `c = -1`.
2. Predict and check the result of each logical operator for `b = [2 31 -40 0]` and `c = 0`.
3. Define two random vectors (`randn(1,7)`) and perform all logical operations, including `xor`, `any` and `all`.

■

Exercise 16.2.

Exercise with logical and relational operators:

1. Let `x = [1 5 2 8 9 0 1]` and `y = [5 2 2 6 0 0 2]`. Execute and explain the results of the following commands:
 - `x > y`
 - `y < x`
 - `x == y`
 - `x <= y`
 - `y >= x`
 - `x | y`
 - `x & (~y)`
 - `(x > y) | (y < x)`
 - `(x > y) & (y < x)`

2. Let $x = 1 : 10$ and $y = [3 \ 5 \ 6 \ 1 \ 8 \ 2 \ 9 \ 4 \ 0 \ 7]$. The exercises here show the techniques of logical-indexing. Execute and interpret the results of the following commands:

- $(x > 3) \& (x < 8)$
- $x((x < 2) \mid (x \geq 8))$
- $x(x > 5)$
- $y((x < 2) \mid (x \geq 8))$
- $y(x \leq 4)$
- $x(y < 0)$

■

Exercise 16.3.

Let $x = [3 \ 16 \ 9 \ 12 \ -1 \ 0 \ -12 \ 9 \ 6 \ 1]$. Provide the command(s) that will:

- set the positive values of x to zero;
- set values that are multiples of 3 to 3 (make use of `rem`);
- multiply the even values of x by 5;
- extract the values of x that are greater than 10 into a vector called y ;
- set the values in x that are less than the mean to 0;
- set the values in x that are above the mean to their difference from the mean.

■

Exercise 16.4.

Execute the following commands and try to understand how z is defined.

```
1 >> hold on
2 >> x = -3:0.05:3; y = sin(3*x);
3 >> subplot(1,2,1); plot(x,y); axis tight
4 >> z = (y < 0.5) .* y;
5 >> subplot(1,2,2); plot(x,y,'r:'); plot(x,z,'r'); axis tight
6 >> hold off
```

■

Before moving on, check whether you now understand the following relations:

```
1 >> a = randperm(10); % random permutation
2 >> b = 1:10;
3 >> b - (a ≤ 7) % subtracts from b a 0–1 vector, taking 1 for
4 >> % a ≤ 7 and 0 otherwise
5 >> (a ≥ 2) & (a < 4) % returns ones at positions where 2 ≤ a < 4
6 >> ¬(b > 4) % returns ones at positions where b ≤ 4
7 >> (a == b) | b == 3 % returns ones at positions where a is equal to b or
8 >> % b is equal to 3
9 >> any(a > 5) % returns 1 when ANY of the a elements are larger than 5
10 >> any(b < 5 & a > 8) % returns 1 when there in the evaluated expression
11 >> % (b < 5 & a > 8) appears at least one 1
12 >> all(b > 2) % returns 1 when ALL b elements are larger than 2
```

16.1.1 The command find

You can extract all elements from the vector or the matrix satisfying a given condition, e.g. equal to 1 or larger than 5, by using logical addressing. The same result can be obtained via the command `find`, which return the positions (indices) of such elements. For instance:


```

1 >> x = [1 1 3 4 1];
2 >> i = (x == 1)
3 i =
4      1      1      0      0      1
5 >> y = x(i)
6 y =
7      1      1      1
8 >> j = find(x == 1)    % j holds indices of those elements satisfying x == 1
9 j =
10     1      2      5
11 >> z = x(j)
12 z =
13     1      1      1

```

Another example is:

```

1 >> x = -1:0.05:1;
2 >> y = sin(x) .* sin(3*pi*x);
3 >> plot (x,y, '-'); hold on
4 >> k = find (y ≤ -0.1)
5 k =
6      9      10      11      12      13      29      30      31      32      33
7 >> plot (x(k), y(k), 'ro');
8 >> r = find (x > 0.5 & y > 0)
9 r =
10     35     36     37     38     39     40     41
11 >> plot (x(r), y(r), 'r*');

```

`find` operates in a similar way on matrices:

```

1 >> A = [1 3 -3 -5; -1 2 -1 0; 3 -7 2 7];
2 >> k = find (A ≥ 2.5)
3 k =
4      3
5      4
6     12
7 >> A(k)
8 ans =
9      3
10     3
11     7

```

In this way, `find` reshapes first the matrix `A` into a column vector, i.e. it operates on `A(:)` in which all columns are concatenated one after another. Therefore, `k` is a list of indices of elements larger than or equal to 2.5 and `A(k)` gives the values of the selected elements. Also the row and column indices can be returned, as shown below:

```

1 >> [I,J] = find (A ≥ 2.5)
2 I =
3      3
4      1
5      3
6 J =
7      1
8      2
9      4
10 >> [A(I(1),J(1)), A(I(2),J(2)), A(I(3),J(3))]    % lists the values
11 ans =

```

12	3	3	7
----	---	---	---

Exercise 16.5.

Let $A = \text{ceil}(5 * \text{randn}(6, 6))$. Perform the following:

- find the indices and list all elements of A which are smaller than -3;
- find the indices and list all elements of A which are smaller than 5 and larger than -1;
- remove those columns of A which contain at least one 0 element.

Exercise with both: logical indexing and the command `find`. ■

16.2 Conditional code execution

16.2.1 Using `if ... elseif ... else ... end`

Selection control structures, `if`-blocks, are used to decide which instruction to execute next depending whether *expression* is TRUE or not. The general description is given below. In the examples below the command `disp` is frequently used. This command displays on the screen the text between the quotes.

- `if ... end`

Syntax

```
if logical_expression
    statement1
    statement2
    ....
end
```

Example

```
if (a > 0)
    b = a;
    disp('a is positive');
end
```

- `if ... else ... end`

Syntax

```
if logical_expression
    block of statements
    evaluated if TRUE
else
    block of statements
    evaluated if FALSE
end
```

Example

```
if (temperature > 100)
    disp('Above boiling. ');
    toohigh = 1;
else
    disp('Temperature is OK. ');
    toohigh = 0;
end
```

- if ... elseif ... else ... end

Syntax

```

if logical_expression1
    block of statements evaluated
    if logical_expression1 is TRUE
elseif logical_expression2
    block of statements evaluated
    if logical_expression2 is TRUE
else
    block of statements evaluated
    if no other expression is TRUE
end

```

Example

```

if (height > 190)
    disp ('very tall');
elseif (height > 170)
    disp ('tall');
elseif (height < 150)
    disp ('small');
else
    disp ('average');
end

```

Exercise 16.6.

In each of the following questions, evaluate the given code fragments. Investigate each of the fragments for the various starting values given on the right. Use MATLAB to check your answers (be careful, since those fragments are not always the proper MATLAB expressions):

- | | |
|------------------------------------------------------------------------------|--------------------------------------------------------------------|
| <pre> 1 if n > 1 2 m = n + 2 3 else 4 m = n - 2 5 end </pre> | <pre> a) n = 7 m = ? b) n = 0 m = ? c) n = -7 m = ? </pre> |
|------------------------------------------------------------------------------|--------------------------------------------------------------------|

- | | |
|-------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------|
| <pre> 1 if s ≤ 1 2 t = 2z 3 elseif s < 10 4 t = 9 - z 5 elseif s < 100 6 t = sqrt(s) 7 else 8 t = s 9 end </pre> | <pre> a) s = 1 t = ? b) s = 7 t = ? c) s = 57 t = ? d) s = 300 t = ? </pre> |
|-------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------|

- | | |
|-----------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------|
| <pre> 1 if t ≥ 24 2 z = 3t + 1 3 elseif t < 9 4 z = t²/3 - 2t 5 else 6 z = -t 7 end </pre> | <pre> a) t = 50 h = ? b) t = 19 h = ? c) t = -6 h = ? d) t = 0 h = ? </pre> |
|-----------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------|

- | | |
|---------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------|
| <pre> 1 if 0 < x < 7 2 y = 4x 3 elseif 7 < x < 55 4 y = -10x 5 else 6 y = 333 7 end </pre> | <pre> a) x = -1 y = ? b) x = 5 y = ? c) x = 30 y = ? d) x = 56 y = ? </pre> |
|---------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------|

■

Exercise 16.7.

Create a script that asks for a number N and computes the drag coefficient C , depending on the Reynold's number N (make use of the `if ... elseif ...` construction). The command `input` might be useful here (use `help` if needed).

$$C = \begin{cases} 0, & N \leq 0 \\ 24/N, & N \in (0, 0.1] \\ 24/N (1 + 0.14 N^{0.7}), & N \in (0.1, 1e3] \\ 0.43, & N \in (1e3, 5e5] \\ 0.19 - 8e4/N, & N > 5e5 \end{cases}$$

Check whether your script works correctly. Compute the drag coefficient for e.g. $N = -3e3, 0.01, 56, 1e3, 3e6$ (remember that e.g. $3e3$ is a MATLAB notation of $3 * 10^3$). ■

Exercise 16.8.

Write a script that asks for an integer and checks whether it can be divided by 2 or 3. Consider all possibilities, such as: divisible by both 2 and 3, divisible by 2 and not by 3 etc (use the command `rem`). ■

16.2.2 Using switch

Another selection structure is `switch`, which switches between several cases depending on an expression, which is either a scalar or a string.

Syntax

```
switch expression
    case choice1
        block of commands1
    case {choice2a, choice2b,...}
        block of commands2
    ...
    otherwise
        block of commands
end
```

Example

```
method = 2;
switch method
    case {1,2}
        disp('Method is linear.');
```

```
    case 3
        disp('Method is cubic.');
```

```
    case 4
        disp('Method is nearest.');
```

```
    otherwise
        disp('Unknown method.');
```

```
end
```

The statements following the first `case` where the expression matches the choice are executed. This construction can be very handy to avoid long `if .. elseif ... else ... end` constructions. The expression can be a scalar or a string. A scalar expression matches a choice if `expression == choice`. A string expression matches a choice if `strcmp(expression, choice)` returns 1 (is true) (`strcmp` compares two strings).

Important: Note that the `switch`-construction only allows the execution of one group of commands.

Exercise 16.9.

Assume that the months are represented by numbers from 1 to 12. Write a script that asks you to provide a month and returns the number of days in that particular month. Alternatively, write a script that asks you to provide a month name (e.g. 'June') instead of a number. Use the `switch`-construction. ■

16.3 Loops

Iteration control structures, *loops*, are used to repeat a block of statements until some condition is satisfied. Two types of loops exist:

- the `for` loop that repeats a group of statements a *fixed* number of times;

Syntax

```
for index = first:step:last
    block of statements
end
```

Example

```
sumx = 0;
for i=1:length(x)
    sumx = sumx + x(i);
end
```

You can specify any `step`, including a negative value. The `index` of the `for`-loop can also be a vector. See some examples of possible variations:

Example 1

```
for i=1:2:n
    ...
end
```

Example 2

```
for i=n:-1:3
    ....
end
```

Example 3

```
for x=0:0.5:4
    disp(x^2);
end
```

Example 4

```
for x=[25 9 81]
    disp(sqrt(x));
end
```

- `while` loop, which evaluates a group of commands as long as *expression* is TRUE.

Syntax

```
while expression
    statement1
    statement2
    statement3
    ...
end
```

Example

```
N = 100;
iter = 1;
msum = 0;
while iter <= N
    msum = msum + iter;
    iter = iter + 1;
end;
```

A simple example how to use the loop construct can be to draw graphs of $f(x) = \cos(nx)$ for $n = 1, \dots, 9$ in different subplots. Execute the following script:

```
1 figure
2 hold on
3 x = linspace(0,2*pi);
4 for n=1:9
```

```

5 subplot(3,3,n);
6 y = cos(n*x);
7 plot(x,y);
8 axis tight
9 end

```

Given two vectors \mathbf{x} and \mathbf{y} , an example use of the loop construction is to create a matrix \mathbf{A} whose elements are defined, e.g. as $A_{ij} = x_i y_j$. Enter the following commands to a script:

```

1 n = length(x);
2 m = length(y);
3 for i=1:n
4     for j=1:m
5         A(i,j) = x(i) * y(j);
6     end
7 end

```

and create \mathbf{A} for $\mathbf{x} = [1 \ 2 \ -1 \ 5 \ -7 \ 2 \ 4]$ and $\mathbf{y} = [3 \ 1 \ -5 \ 7]$. Note that \mathbf{A} is of size n -by- m . The same problem can be solved by using the **while**-loop, as follows:

```

1 n = length(x);
2 m = length(y);
3 i = 1; j = 1; % initialize i and j
4 while i ≤ n
5     while j ≤ m
6         A(i,j) = x(i) * y(j);
7         j = j+1; % increment j; it does not happen automatically
8     end
9     i = i+1; % increment i
10 end

```

Exercise 16.10.

Determine the sum of the first 50 squared numbers with a control loop. ■

Exercise 16.11.

Write a script to find the largest value n such that the sum: $\sqrt{1^3} + \sqrt{2^3} + \dots + \sqrt{n^3}$ is less than 1000. ■

Exercise 16.12.

Use a loop construction to carry out the computations. Write short scripts.

- Given the vector $\mathbf{x} = [1 \ 8 \ 3 \ 9 \ 0 \ 1]$, create a short set of commands that will:
 - add up the values of the elements (check with **sum**);
 - computes the running sum (for element j , the running sum is the sum of the elements from 1 to j ; check with **cumsum**);
 - computes the sine of the given x -values (should be a vector).
- Given $\mathbf{x} = [4 \ 1 \ 6 \ -1 \ -2 \ 2]$ and $\mathbf{y} = [6 \ 2 \ -7 \ 1 \ 5 \ -1]$, compute matrices whose elements are created according to the following formulas:
 - $a_{ij} = y_i/x_j$;
 - $b_i = x_i y_i$ and add the elements in **btot**;
 - $c_{ij} = x_i/(2 + x_i + y_j)$;
 - $d_{ij} = 1/\max(x_i, y_j)$.

3. Write a script that transposes a matrix **A**. Check its correctness with the MATLAB operation: **A'**.
4. Create an m -by- n array of random numbers (use **rand**). Move through the array, element by element, and set any value that is less than 0.5 to 0 and any value that is greater than (or equal to) 0.5 to 1.
5. Write a script that will use the random-number generator **rand** to determine:
 - the number of random numbers it takes to add up to 10 (or more);
 - the number of random numbers it takes before a number between 0.8 and 0.85 occurs;
 - the number of random numbers it takes before the mean of those numbers is within 0.01 and 0.5.

It will be worthwhile to run your script several times because you are dealing with random numbers. Can you predict any of the results that are described above?



Exercise 16.13.

Write a script that asks for a temperature in degrees Celsius **tc** and computes the equivalent temperature in degrees Fahrenheit **tf** (use the formula $tf = 9/5 * tc + 32$). The script should keep running until no number is provided to convert. The functions **input** and **isempty** (use **help** to learn more) should be useful here. ■

16.4 Evaluation of logical and relational expressions in the control flow structures

The relational and logical expressions may become more complicated. It is not difficult to operate on them if you understand how they are evaluated. To explain more details, let us consider the following example:

```
1 if (~isempty(data)) & (max(data) < 5)
2     ....
3 end
```

This construction of the **if**-block is necessary to avoid comparison if **data** happens to be an empty matrix. In such a case you cannot evaluate the right logical expression and MATLAB gives an error. The **&** operator returns 1 only if both expressions: **~isempty (data)** and **max(data) < 5** are true, and 0 otherwise. When **data** is an empty matrix, the next expression is not evaluated since the whole **&**-expression is already known to be false. The second expression is checked only if **data** is a non-empty matrix. Remember to put logical expression units between brackets to avoid wrong evaluations!

Important: The fact that computers make use of floating-point arithmetic means that often you should be careful when comparing two floating-point numbers just by:

```
1 if (x == y)
2     ....
3 end
```

(Of course, such a construction is allowed e.g. when you know that x and y represent integers.) Instead of the above construction, you may try using this:

```
1 if (abs (x - y) < tolerance)           % e.g. tolerance = 1e-10
2     ....
3 end
```

Exercise 16.14.

Consider the following example:

```
1 max_iter = 50;
2 tolerance = 1e-4;
3 iter = 0;
4 xold = 0.1;
5 x = 1;
6 while (abs (x - xold) > tolerance) & (iter < max_iter)
7     xold = x;
8     x = cos(xold);
9     iter = iter + 1;
10 end
```

This short program tries to solve the equation $\cos(x) = x$ (x is the solution found). Make the script `solve_cos` from the presented code. Note that the `while`-loop is repeated as long as both conditions are true. If either the condition $(|x - xold| \leq \text{tolerance})$ or $(\text{iter} \geq \text{max_iter})$ is fulfilled then the loop is quitted. Run the script `solve_cos` for different `tolerance` parameters, e.g.: $1e-3$, $1e-6$, $1e-8$, $1e-10$ etc. Use `format long` to check more precisely how much the found x is really different from $\cos(x)$. For each `tolerance` value check the number of performed iterations (`iter` value). ■

Exercise 16.15.

Create the script `solve_cos2`, which is equal to the one given, replacing the `while`-loop condition by:

```
1 while (abs (x - xold) > tolerance) | (iter < max_iter)
```

Try to understand the difference and confirm your expectations by running `solve_cos2`. What happens to `iter`? ■

16.5 Exercises

Exercise 16.16.

The circumference of a unit circle (with radius 1) equals 2π . Knowing this, we can approximate π or 2π as follows. Section the circle into triangles, figure 16.1. Starting with 4 triangles the length of s_n equals $\sqrt{2}$. A first approximation for the circumference of the circle is $4\sqrt{2} \approx 5.66$.

Next the number of triangles is doubled. Length a is half of the previously determined length s_n . Then subsequently lengths b and c can be computed and also the length of the new section s_{n+1} is known from which a new (and improved) approximation for the circumference of the circle can be obtained.




Figure 16.1: Sectioning a circle into triangle to approximate π . From the BEAM-B handout.

Write a script to evaluate the approximation for the circumference of the circle while the number of triangles is doubled at least 20 times.

■

Chapter 17

Functions

M-files have been introduced in chapter 14 to collect MATLAB commands in *scripts*. In this chapter M-files are used to define MATLAB *functions*.

17.1 Function m-file

Functions m-files are true subprograms, since they take input arguments and/or return output parameters. They can call other functions, as well. Variables defined and used inside a function, different from the input/output arguments, are invisible to other functions and the command environment. The general syntax of a function is:

```
1 function [outputArgs] = function_name (inputArgs)
```

outputArgs are enclosed in []:

- a comma-separated list of variable names;
- [] is optional when only one argument is present;
- functions without **outputArgs** are legal¹.

inputArgs are enclosed in ():

- a comma-separated list of variable names;
- functions without **inputArgs** are legal.

MATLAB provides a structure for creating your own functions. The first line of the file should be a definition of a new function (also called a header). After that, a continuous sequence of comment lines should appear. Their goal is to explain what the function does, especially when this is not trivial. Not only a general description, but also the expected input parameters,

¹In other programming languages, functions without output arguments are called *procedures*

returned output parameters and synopsis should appear there. The comment lines (counted up to the first non-comment line) are important since they are displayed in response to the `help` command. Finally, the remainder of the function is called *the body*. Function m-files terminate execution and return when they reached the end of the file or, alternatively, when the command `return` is encountered. As an example, the function `average` is defined as follows:

The diagram shows the following code with annotations:

```

function avr = average (x)
% AVERAGE computes the average value of a vector x
%           and returns it in avr
% Notes: an example of a function
n = length(x);
avr = sum(x)/n;
return;

```

- the first line must be the function definition**: points to `function avr = average (x)`
- output argument**: points to `avr`
- function name**: points to `average`
- input argument**: points to `(x)`
- comment**: points to the first two comment lines.
- function body**: points to the code lines `n = length(x);`, `avr = sum(x)/n;`, and `return;`.
- a blank line within the comment; Notes information will NOT appear when you ask: help average**: points to the blank line between the first two comment lines.

Important: The name of the function and the name of the file stored on disk should be *identical*. In our case, the function should be stored in a file called `average.m`.

Exercise 17.1.

Create the function `average` and store it on disk as `average.m`. Remember about the comment lines. Check its usability by calling `help average`. Run `average.m` by typing `avr1 = average(1:10);` ■

Here is an another example of a function:

```

1 function [avr,sd] = stat(x)
2 %STAT Simple statistics.
3 %   Computes the average value and the standard deviation of a vector x.
4 n = length(x);
5 avr = sum(x)/n;
6 sd = sqrt(sum((x - avr).^2)/n);
7 return;

```

Warning: The functions `mean` and `std` already exist in MATLAB. As long as a function name is used as variable name, MATLAB can not perform the function. Many other, easily appealing names, such as `sum` or `prod` are reserved by MATLAB functions, so be careful when choosing your names (see section 18.3).

The `return` statement can be used to force an early return. An exemplary use of the `return` is given below:

```

1 function d = determinant(A)
2 %DETERMINANT Computes the determinant of a matrix A
3 [m,n] = size(A);
4 if (m ≠ n)
5     disp ('Error. Matrix should be square. ');
6     return;
7 else
8     d = det(A); % standard Matlab function
9 end

```

```
10 return;
```

The use of the `return` command in combination with a check of the *number* of input arguments is given in the function `checkarg`, which is presented in section 17.1.2.

When controlling the proper use of parameters, the function `error` may prove useful. It displays an error message, aborts function execution, and returns to the command environment. Here is an example:

```
1 if (a ≥ 1)
2     error ('a must be smaller than 1');
3 end
```

Exercise 17.2.

Change some scripts that you created into functions, e.g. create the function `drag`, computing the drag coefficient (see section 16.2), or `solve_cos` (see section 16.4) or `cubic_roots` (see section 14). ■

Exercise 17.3.

Write the function `[elems, mns] = nonzero(A)` that takes as the input argument a matrix `A` and returns all nonzero elements of `A` in the column vector `elems`. The output parameter `mns` holds values of the means of all columns of `A`. ■

Exercise 17.4.

Write the function `[meanrow, meancol] = allmeans(A)` that takes as the input argument a matrix `A` and returns the means of all its rows and columns respectively. ■

Exercise 17.5.

Create the function `[A,B,C] = sides(a,b,c)` that takes three positive numbers `a`, `b` and `c`. If they are sides of a triangle, then the function returns its angles `A`, `B` and `C`, measured in degrees. Display an error when necessary. ■

Exercise 17.6.

The area of a triangle with sides of length a , b , and c is given by: $ar = \sqrt{s(s-a)(s-b)(s-c)}$, where $s = (a+b+c)/2$. Write a function that accepts a , b and c as inputs and returns the value ar as output. Note that the sides should be non-negative and should fulfill the triangle inequality. Make use of the `error` command. ■

Exercise 17.7.

Create the function `randint` that randomly generates a matrix of integer numbers (use the command `rand`). These integers should come from the interval $[a, b]$. Exercise in documenting the function. Use the following function header:

```
1 function r = randint(m,n,a,b) % a m-by-n matrix
```

If this seems too difficult, start first with the fixed interval, e.g. $[a, b] = [0, 5]$ (and remove `a` and `b` from the function definition) and then try to make it work for an arbitrary interval. ■

17.1.1 Subfunctions

A function m-file may contain more than a single function. The function appearing first in the m-file shares the file name, as mentioned before. Other functions are *subfunctions* and can be called from the primary function only. So, they cannot be accessed from outside the file, in which they are defined (neither from the Command Window nor via other m-files). A subfunction is created by defining a new function with the **function** statement after the body of the preceding function. The use of subfunctions is recommended to keep the function readable when it becomes too long and too complicated. For example, **average** is now a subfunction within the file **stat.m**:

```
1 function [a,sd] = stat(x)
2 %STAT Simple statistics.
3 %     Computes the average value and the standard deviation of a vector x.
4 n = length(x);
5 a = average(x,n);
6 sd = sqrt(sum((x - avr).^2)/n);
7 return;
8
9 function a = average (x,n)
10 %AVERAGE subfunction
11 a = sum(x)/n;
12 return;
```

Exercise 17.8.

Modify the function **stat** presented above by using the subfunction **average** for the computation of variable **sd** as well. ■

17.1.2 Special function variables

Each function has two internal variables: **nargin** - the number of function input arguments that were used to call the function and **nargout** - the number of output arguments. Analyze the following function:

```
1 function [out1,out2] = checkarg (in1,in2,in3)
2 %CHECKARG Demo on using the nargin and nargout variables.
3 if (nargin == 0)
4     disp('no input arguments');
5     return;
6 elseif (nargin == 1)
7     s = in1;
8     p = in1;
9     disp('1 input argument');
10 elseif (nargin == 2)
11     s = in1+in2;
12     p = in1*in2;
13     disp('2 input arguments');
14 elseif (nargin == 3)
15     s = in1+in2+in3;
16     p = in1*in2*in3;
17     disp('3 input arguments');
18 else
19     error('Too many inputs.');
```

```

20 end
21
22 if (nargout == 0)
23     return;
24 elseif (nargout == 1)
25     out1 = s;
26 else
27     out1 = s;
28     out2 = p;
29 end

```

Exercise 17.9.

Construct the function `checkarg`, call it with different number of input and output arguments and try to understand its behavior. For example:

```


1 >> checkarg
2 >> s = checkarg(-6)
3 >> s = checkarg(23,7)
4 >> [s,p] = checkarg(3,4,5)

```

■

17.1.3 Local and global variables

Each m-file function has access to a part of memory separate from MATLAB's workspace. This is called the *function workspace*. This means that each m-file function has its own **local** variables, which are separate from those of other function and from the workspace. To understand it better analyze the following diagram:



In the MATLAB workspace, variables `a`, `b` and `c` are available. Variables `x`, `y` and `z` are visible only in the function `myfun`. However, if several functions and/or the workspace, **all** declare a particular variable as **global**, then they all share this variable (see `help global`). Any assignment to that variable is available to all other functions and/or the workspace. However, you should be careful when using global variables. It is very easy to get confused and end up with serious errors.

OPTIONAL

17.1.4 Indirect function evaluation

Using indirect function evaluation makes programming even more general, since functions can become input arguments. The crucial MATLAB command here is `feval`, an abbreviation of

function **evaluation**. The **feval** command allows execution of a function specified by a string. The general definition is as follows:

```
1 [y1,...,yn] = feval (F,x1,...,xn),
```

where **F** is a name of a function defined in MATLAB, **x1,...,xn** are input arguments and **y1,...,yn** are possible output parameters. Consider an example:

```
1 >> x = pi; y = cos(x);
2 >> z = feval('cos',x);
```

The last command is also equivalent to the following two expressions:

```
1 >> F = 'cos';
2 >> z = feval(F,x)
```

Indirect function evaluation is a nice tool to build a program with a function given as an argument.

Exercise 17.10.

Create the function **funplot** and try to understand how it works:

```
1 function funplot (F, xstart, xend, col);
2 %FUNPLOT makes a plot of the function F at the interval [xstart, xend].
3 %       The plot should be made in one of the standard Matlab colors, so
4 %       'col' is one of the following value: 'b','k','m','g','w','y' or 'r'.
5 % default values:
6 %       [xstart,xend] = [0,10]
7 %       col = 'b'
8
9 % Note: illustrates the use of feval command
10
11 if (nargin == 0)
12     error ('No function is provided.');

```

```

36     otherwise
37         error ('Wrong col value provided.')
38     end
39
40     x = linspace(xstart, xend);
41     y = feval(F,x);
42     plot (x,y,col);
43     description = ['Plot of ', F];
44     title (description);
45     return;

```

Note the use of comments, the `nargin` variable and the `switch`-construction. Call `funplot` for different built-in functions, like `sin`, `exp`, etc. Test it for your own functions as well. Write for example a function `myfun` that computes $\sin(x \cos(x))$ or $\log(|x \sin(x)|)$. Explain why it would be wrong to use the fragment given below instead of its equivalent part in `funplot`.

```

1  if nargin < 2
2      xstart = 0;
3      xend   = 10;
4  elseif nargin < 3
5      error ('Wrong number of arguments. You should provide xstart and xend.');
```

■

END OPTIONAL

17.2 Scripts vs. functions

The most important difference between a script and a function is that *all* script's parameters and variables are externally accessible (i.e. in the workspace), where function variables are not. Therefore, a script is a good tool for documenting work, designing experiments and testing. In general, create a function to solve a given problem for arbitrary parameters. Use a script to run functions for specific parameters required by the assignment.

17.3 Exercises

Exercise 17.11.

Create a MATLAB function to evaluate the expression

$$f(n) = \left(1 + \frac{1}{n}\right)^n$$

as a function of n . Next check the outcome of this expression for increasing n . Can you verify that $\lim_{n \rightarrow \infty} f(n) = e$? ■

Exercise 17.12.

Create the function `binom` that computes the value of the binomial symbol $\binom{n}{k}$. Make the

function header:

function **b** = **binom** (**n,k**). Note that in order to write this function, you will have to create the **factorial** function, which computes the value of $n! = 1 * 2 * \dots * n$. This may become a separate function (enclosed in a new file) or a subfunction in the **binom.m** file. Try to implement both cases if you got acquainted with the notion of a subfunction. Having created the **binom** function, write a script that displays on screen all the binomial symbols for $n = 8$ and $k = 1, 2, \dots, 8$ or write a script that displays on screen the following 'triangle' (use the **fprintf** command; try **help fprintf**. Note **fprintf** is a c-command, which is in the header **cstdio** (with the same syntax), but **c++** streams are much better for dealing with input/output, so its use is not recommended in **c++**.

$$\begin{array}{cccc} \binom{1}{1} & & & \\ \binom{2}{1} & \binom{2}{2} & & \\ \binom{3}{1} & \binom{3}{2} & \binom{3}{3} & \\ \binom{4}{1} & \binom{4}{2} & \binom{4}{3} & \binom{4}{4} \end{array}$$

■

Exercise 17.13.

Solve the following exercises by using either a script or a function. The nature of the input/output and display options is left to you. Problems are presented with the increasing difficulty; start simple and add complexity. If possible, try to solve all of them.

1. Write a function which computes the cumulative product of a vector elements. The cumulative product up to x_j - the j -th element of the vector x is defined by $C_j = \prod_{k=1}^j x_k$ for $j = 1 : \text{length}(x)$. Check your results with the built-in function **cumprod**.
2. Write MATLAB function **m** = **wmean** (**x,w**) computing the weighted arithmetic mean, given the vector **x** and a vector of nonnegative weights **w**, such that $\sum_{i=1}^n w_i > 0$, i.e. $\frac{\sum_{i=1}^n w_i x_i}{\sum_{i=1}^n w_i}$. Add error messages to terminate execution of the function in the case when:
 - **x** and **w** are of different lengths,
 - at least one element of **w** is negative,
 - sum of all weights is equal to zero.
3. What is the greatest value of n that can be used in the sum

$$1^2 + 2^2 + 3^2 + \dots + n^2$$

and gives a total value of the sum less than 100?

4. Compute the value of π using the following series:

$$\frac{\pi^2 - 8}{16} = \sum_{n=1}^{\infty} \frac{1}{(2n-1)^2 (2n+1)^2}$$

How many terms are needed to obtain an accuracy of $1e-12$? How accurate is the sum of 100 terms?

5. Write a program that approximates π by computing the sum:

$$\frac{\pi}{4} \approx \sum_{n=0}^m \frac{(-1)^n}{2n+1}$$

The more terms in summation the larger the accuracy (although this is not an efficient formula, since you add and subtract numbers). How many terms are needed to approximate π with 5 decimals? Use the sum to approximate π using 10, 100, 1e3, 1e4, 5e4, 1e5, 5e5 and 1e6 terms. For each of these numbers compute the approximation error. Plot the error as a function of the term numbers used in a summation.

6. The Fibonacci numbers are computed according to the following relation:

$$F_n = F_{n-1} + F_{n-2}, \quad \text{with } F_0 = F_1 = 1$$

- Compute the first 10 Fibonacci numbers.
 - For the first 50 Fibonacci numbers, compute the ratio $\frac{F_n}{F_{n-1}}$. It is claimed that this ratio approaches the value of the golden mean $\frac{1+\sqrt{5}}{2}$. What do your results show?
7. Consider a problem of computing the n -th Fibonacci number. Find three different ways to implement this and construct three different functions, say `fib1`, `fib2` and `fib3`. Measure the execution time of each function (use the commands `tic` and `toc`) for, say, $n = 20, 40$ or 100.
8. The Legendre polynomials $P_n(x)$ are defined by the following recurrence relation:

$$(n+1)P_{n+1}(x) - (2n+1)xP_n(x) + nP_{n-1}(x) = 0$$

with $P_0(x) = 1$, $P_1(x) = x$ and $P_2(x) = (3x^2 - 1)/2$. Compute the next three Legendre polynomials and plot all 6 over the interval $[-1, 1]$.

9. Write a script that asks for an integer n (or a function that has n as the input argument) and then computes the following: while the value of n is greater than 1, replace the integer with $(n/2)$ if the integer is even. Otherwise, replace the integer with $(3n + 1)$. Make provision to count the number of values in (or the length of) the sequence that results.

Example calculation: If $n = 10$, the sequence of integers is 5, 16, 8, 4, 2, 1, so the length is 6.

Make a plot of the length of the sequence that occurs as a function of the integers from 2 to 30. For example, when $n = 10$, the length is 6 while for $n = 15$, the length is 17. Is there any pattern? Try larger numbers to see if any pattern occurs. Is there any integer for which the sequence does not terminate?

10. Provide all prime numbers smaller than the given N .

■

Exercise 17.14.

Declare an array with 10 001 entries. Compute the solution vector $f(x_i)$, with $x_i = 0, \dots, 10\,000$

$$f(x_i) = \frac{x_i}{\sin(x_i) + 2}$$

In the same program, form the alternating sum

$$S_n = f(x_0) + f(x_1) - f(x_2) + f(x_3) - \dots$$

up to the last term and print the result to the screen.

■

Chapter 18

Writing and debugging MATLAB programs

The recommendations in this section are general for programming in any language. Learning them now will turn out to be beneficial in the future or while learning real programming languages like *C/C++*, where structured programming is indispensable.

18.1 Structural programming

Never write all code at once; program in small steps and make sure that each of these small steps works as expected, before proceeding to the next one. Each step should be devoted to only one task. Do not solve too many tasks in one module, because you may easily make a mess. This is called a *structured* or *modular* way of programming. Formally, modularity is the hierarchical organization of a system or a task into self-contained subtasks and subsystems, each having a prescribed input-output communication. It is an essential feature of a well designed program. The benefit of structural programming are: easier error detection and correction, modifiability, extensibility and portability. A general approach to a program development is presented below:

1. Specification.

Read and understand the problem. The computer cannot do anything itself: you have to tell it how to operate. Before using the computer, some level of preparation and thought is required. Some basic questions to be asked are:

- What are the parameters/inputs for this problem?
- What are the results/outputs for this problem?
- What form should the inputs/outputs be provided in?
- What sort of algorithms is needed to find the outputs from the inputs?

2. Design.

Split your problem into a number of smaller and easier tasks. Decide how to implement them. Start with a schematic implementation to solve your problem, e.g. create function

headers or script descriptions (decide about the input and output arguments). To do this, you may use, for example, a top-down approach. You start at the most general level, where your first functions are defined. Each function may be again composed of a number of functions (subfunctions). While 'going down' your functions become more precise and more detailed.

As an example, imagine that you have to compare the results of the given problem for two different datasets, stored in the files `data1.dat` and `data2.dat`. Schematically, such a top-down approach could be designed as:

◊ This is the top (the most general) level. A script `solve_it` is created:

```
1 [d1, d2] = read_data ('data1.dat', 'data2.dat');
2 [res1, err1] = solve_problem (d1);
3 [res2, err2] = solve_problem (d2);
4 compare_results (res1, res2, err1, err2);
```

◊ This is the second level. The functions `read_data`, `solve_problem` and `compare_results` belong here. Each of them has to be defined in a separate file:

- ```
1 function [d1, d2] = read_data (fname1, fname2)
2 % Here should be some description.
3 %
4 fid1 = fopen (fname1, 'w');
5 % check whether the file fname1 exists
6 fclose(fid1);
7 fid2 = fopen (fname2, 'w');
8 % check whether the file fname2 exists
9 fclose(fid2);
10
11 d1 = ...
12 d2 = ...
13 return;
```

- ```
1 function [res, err] = solve_problem (d)
2 % Here should be some (possibly detailed) description.
3 %
4 ....
5 res = ... % the data d is used to compute res
6 err = compute_error (res);
7 return;
```

- ```
1 function compare_results (res1, res2, err1, err2)
2 % Some description.
3 tol = 1e-6;
4
5 if abs (err1 - err2) > tol
6 fprintf ('The difference is significant.')
7 else
8 fprintf ('The difference is NOT significant.')
9 end;
10 return;
```

◊ In this example, this is the last level. The function `solve_problem` uses the function: `compute_error`, which has to be defined:

- ```
1 function err = compute_error (res)
2 % Here should be some (possibly detailed) description.
3 %
```

```

4  ....
5  err = ....
6  return;

```

% the variable res is used to compute err

3. Coding.

Implement the algorithms sequentially (one by one). Turning your algorithm into an efficient code is not a one-shot process. You will have to try, make errors, correct them and even modify the algorithm. So, **be patient**. While implementing, make sure that all your outputs are computed at some point. Remember about the comments and the style requirements (see section 18.3).

4. Running and debugging (see also section 18.2).

Bugs will often exist in a newly written program. Never, ever, believe or assume that the code you just created, works. **Always check the correctness of each function or script: Twice.** You may add some extra lines to your code which will present the intermediate results (screen displays, plots, writes to files) to help you controlling what is going on. Those lines can be removed later.

5. Testing and Verification.

After the debugging process, the testing stage starts. Prepare a number of tests to verify whether your program does what it is expected to do. Remember that good tests are those for which the answers are known. Your program should produce correct results for normal test conditions as well as boundary conditions.

6. Maintenance.

In solving your task, new ideas or problems may appear. Some can be interesting and creative and some can help you to understand the original problem better; you may see an extent to your problem or a way to incorporate new things. If you have a well-designed problem, you will be able to easily modify it after some time. Take a responsibility to improve your solutions or correct your errors when found later.

18.2 Debugging

Debugging is the process by which you isolate and fix any problem with your code. Two kinds of errors may occur: *syntax error* and *runtime error*. Syntax errors can usually be easily corrected by MATLAB error messages. Already while editing a script or function m-file, the MATLAB editor will indicate possible errors. Add e.g. the following line to one of your m-files:

```
1 x = 2pi;
```

There should be a marker on the right side of the edit window. Moving the cursor over this mark will reveal the message with the syntax error: “Parse error at 'pi': usage appears to be invalid MATLAB syntax.”.

Runtime errors are algorithmic in nature and they occur when e.g. you perform a calculation incorrectly. They are usually difficult to track down, but they are apparent when you notice unexpected results.

Debugging is an inevitable process. The best way to reduce the possibility of making a runtime error is *defensive programming*:

- Do not assume that input is correct, simply check.
- Where reasonable and possible, provide a default option or value.
- Provide diagnostic error messages.
- Optionally print intermediate results to check the correctness of your code.

Defensive programming is a part of the early debugging process. Another important part is modularity, breaking large task into small subtasks, which allows for developing tests for each of them more easily. You should always remember to run the tests again after the changes have been made. To make this easy, provide extra print statements that can be turned on or off.

MATLAB provides an interactive debugger. It allows you to set and clear *breakpoints*, specific lines in an m-file at which the execution halts. It also allows you to change the workspace and execute the lines in an m-file one by one. The MATLAB m-file editor also has a debugger. The debugging process can be also done from the command line. To use the debugging facility to find out what is happening, you start with the `dbstop` command. This command provides a number of options for stopping execution of a function. A particularly useful option is:

Next year add some information on the actually debugger, and commands with keyboard etc...

```
1 dbstop if error
```

This stops any function causing an error. Then just run the MATLAB function. Execution will stop at the point where the error occurs, and you will get the MATLAB prompt back so that you can examine variables or step through execution from that point. The command `dbstep` allows you to step through execution one line at a time. You can continue execution with the `dbcont`. To exit debug mode, type `dbquit`. For more information, use `help` for the following topics: `dbstop`, `dbclear`, `dbcont`, `dbstep`, `dbtype`, `dbup` and `dbquit`.

18.3 Recommended programming style

Programming style is a set of conventions that programmers follow to standardize their code to some degree and to make the overall program easier to read and to debug. This will also allow you to quickly understand what you did in your program when you look at it weeks or months from now. The style conventions are for the reader only, but **you** will become that reader one day.

Some style requirements and style guidelines are presented below. These are recommendations, and some personal variations in style are acceptable, but you should not ignore them. It is important to organize your programs properly since it will improve the readability, make the debugging task easier and save time of the potential readers.

1. You should **always** comment difficult parts of the program! But ... do not explain the obvious.
2. Comments describing tricky parts of the code, assumptions, or design decisions are suggested to be placed above the part of the code you are attempting to document. Try to avoid big blocks of comments except for the description of the m-file header.

3. Indent a few spaces (preferably 2 or 3) before lines of the code and comments inside the control flow structures. The layout should reflect the program 'flow'. Here is an example:

```

1 x = 0:0.1:500;
2 for i=1:length(x)
3     if x(i) > 0
4         s(i) = sqrt(x(i));
5     else
6         s(i) = 0;
7     end
8 end

```

The MATLAB m-file editor will introduce such spaces by default.

4. Avoid the use of magic numbers; use a *constant* variable instead. When you need to change the number, you will have to do it only once, rather than searching all over your code. An example:

<pre> 1 % A BAD code that uses 2 % magic numbers 3 r = rand(1,50); 4 for i = 1:50 5 data(i) = i * r(i); 6 end 7 y = sum(data)/50; 8 disp(['Number of points is 50.']); 9 </pre>	<pre> % This is the way it SHOULD be n = 50; % number of points r = rand(1,n); data = (1:n) .* r; avr = sum(data)/n; disp(['Number of points is ',... int2str(n)]); </pre>
-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

5. Avoid the use of more than *one* code statement per line in your script or function m-files.
6. No line of code should exceed 80 characters (it is a rare case when this is not possible). When necessary, use the *Line Continuation Symbol*, i.e. the characters `...`, at the end of a line of to indicate that the expression continues on the next line.
7. Avoid declaring global variables. You will hardly ever encounter a circumstance under which you will really need them. Global variables can get you into trouble without your noticing it!
8. Variables should have meaningful names. You may also use the standard notation, e.g. `x`, `y` are real-valued, `i`, `j`, `k` are indices or `n` is an integer. This will reduce the number of comments and make your code easier to read. However, here are some pitfalls when choosing variable names:
 - A meaningful variable name is good, but when it gets longer than 15 characters, it tends to obscure rather than improve the code readability.
 - Be careful with names since there might be a conflict with MATLAB's built-in functions, or reserved names like *mean*, *end*, *sum* etc (check in index or ask **which <name>** in MATLAB - if you get the response **<name> not found** it means that you can safely use it). Note that even the obvious examples of variables `i` or `j` overwrite the built-in complex imaginary unit *i*.
 - Avoid names that look similar or differ only slightly from each other.
9. Use white spaces; both horizontally and vertically, since it will greatly improve the readability of your program. Blank lines should separate larger blocks of the code.
10. Test your program before submitting it. Do not just assume it works.

Chapter 19

Cell arrays and structures

19.1 Cell arrays

Cell arrays are arrays whose elements are *cells*. Each cell can contain any data, including numeric arrays, strings, cell arrays etc. For instance, one cell can contain a string, another a numeric array etc. Below, there is a schematic representation of an exemplar cell array:

cell 1,1 <div><div>3 1 -7 7 2 4 0 -1 6 7 3 7</div></div>	cell 1,2 <div><div>John Smith</div><div>35</div><div>5900 6000 6100</div></div>	cell 1,3 <div><div>1 + 3i</div></div>
cell 2,1 <div><div>'this is a text'</div></div>	cell 2,2 <div><div><div>'Hi'</div><div>[7,7]</div></div><div><div><div>-1 0 0 -1</div></div><div>'Bye'</div></div></div>	cell 2,3 <div><div><div>'Living'</div><div>'implies effort'</div></div></div>

Cell arrays can be built up by assigning data to each cell. The cell contents are accessed by brackets `{}`. For example:

```
1 >> A(1,1) = {[3 1 -7;7 2 4;0 -1 6;7 3 7]};
2 >> A(2,1) = {'this is a text'};
3 >> B(1,1) = {'John Smith'}; B(2,1) = {35}; B(3,1) = {[5900 6000 6100]}
4 >> A(1,2) = {B}; % cell array B becomes one of the cells of A
5 >> C(1,1) = {'Hi'}; C(2,1) = {[ -1 0;0 -1]}; C(1,2) = {[7,7]}; C(2,2) = {'Bye'};
6 >> A(2,2) = {C}; % cell array C becomes one of the cells of A
7 >> A % A represents now the first 2 columns of the exemplar
8 % cell array
```



```

9      A =
10      [4x3 double]      {3x1 cell}
11      'this is a text'  {2x2 cell}
12  >> A(2,2)             % access the cell but not its contents
13  ans =
14      {2x2 cell}
15  >> A{2,2}             % use {} to display the contents
16  ans =
17      'Hi'              [1x2 double]
18      [2x2 double]      'Bye'
19  >> A{2,2}{2,1}        % take the (2,1) element of the cell A{2,2}
20  ans =
21      -1      0
22      0      -1

```

There are also two useful functions with meaningful names: `celldisp` and `cellplot`. Use `help` to learn more.

The common application of cell arrays is the creation of text arrays. Consider the following example:

```

1  >> M = {'January'; 'February'; 'March'; 'April'; 'May'; 'June'; 'July'; 'August';
2  'September'; 'October'; 'November'; 'December'};
3  >> fprintf('It is %s.\n', M{9});
4  It is September.

```

Exercise 19.1.

Exercise with the concept of a cell array, first by typing the examples presented above. Next, create a cell array `W` with the names of week days, and using the command `fprintf`, display on screen the current date with the day of the week. The goal of this exercise is also to use `fprintf` with a format for a day name and a date, in the spirit of the above example. ■

19.2 Structures

Structures are MATLAB arrays with data objects composed of *fields* and are very similar to classes in `c++`. Each field contains one item of information. For example, one field might include a string representing a name, another a scalar representing age or an array of the last few salaries. Structures are especially useful for creating and handling a database. One of the possible ways to create a structure is by assigning data to individual fields. Imagine that you want to construct a database with some information on workers of a company:

```

1  >> worker.name = 'John Smith';
2  >> worker.age = 35;
3  >> worker.salary = [5900, 6000, 6100];
4  >> worker =
5      name: 'John Smith'
6      age: 35
7      salary: [5900 6000 6100]

```

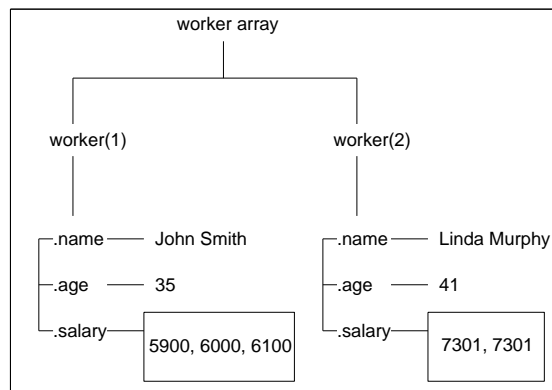
In this way, a 1×1 structure array `worker` with three fields: `name`, `age` and `salary` is constructed. To expand the structure, add a subscript after the structure name:

```

1 >> worker(2).name = 'Linda Murphy'; % after this, a 2nd subarray is created
2 >> worker(2).age = 41;
3 >> worker(2).salary = [7301, 7301]; % field sizes do not need to match!
4 >> worker
5 1x2 struct array with fields:
6     name
7     age
8     salary

```

Since this structure has now the size of 1×2 , MATLAB does not display the contents of all fields. The data are now organized as follows:



Structures can also be build by using the **struct** function. For example:

```

1 >> employee=struct( 'name','John Smith','age',35,'salary',[5900, 6000, 6100]);

```

To create an empty structure with fields field1, field2

```

1 s = struct('field1', {}, 'field2', {})

```

To create an empty structure with no fields

```

1 struct([])

```

To access an entire field, include a field name after a period. To access a subarray, follow the standard way of using subscripts:

```

1 >> worker(1).age
2     35
3 >> worker(2).salary(2)
4     7301
5 >> worker(2)
6     name: 'Linda Murphy'
7     age: 41
8     salary: [7301 7301]

```

An alternative way is to use the **getfield** function:

```

1 >> getfield(worker,{2},'salary')
2          7301          7301

```

There exists also a function `setfield`, which assigns values to a given field. New fields can be added or deleted from every structure. To delete a field, use the command `rmfield`. Analyze the following example:

```

1 >> worker2 = worker;
2 >> worker2 = rmfield (worker, 'age');
3 >> worker2(2).street = 'Bakerstreet 5';
4 >> worker2(2)
5     name: 'Linda Murphy'
6     salary: [7301 7301]
7     street: 'Bakerstreet 5'
8 >> worker2(1)
9     name: 'John Smith'
10    salary: [5900 6000 6100]    % in all other substructures address field is empty
11    street: []                % you should assign new values for each substructure

```

Operating on fields and field elements is done in the same way as on any other array. Consider the following example, where the average salary for each worker is computed.

```

1 avr_salary(1) = mean (worker(1).salary);
2 avr_salary(2) = mean (worker(2).salary);

```

Remark: structures as a concept are common organizations of data in other programming languages. They are created in different ways, but the intention remains the same.

Exercise 19.2.

Construct a structure `friend`, with the following fields: `name`, `address`, `age`, `birthday`. Insert a few friends with related information. Remove e.g. the field `age` and add the field `phone`. ■

Concerning the use of lists and parentheses in MATLAB, please see `help lists` and `help paren`.

General remark: *Classes* and *objects* allow for adding new data types and new operations to MATLAB. For instance, the class of a variable describes the structure of the variables and the operations permitted as well as functions to be used. An object is an instance of a particular class. The use of classes and objects is the base of *object-oriented programming*, which is also possible in MATLAB.

Chapter 20

File input/output operations

MATLAB File input and output (I/O) functions read and write arbitrary binary and formatted text files. This enables you to read data collected in other formats and to save data for other programs, as well. Before reading or writing a file you must open it with the `fopen` command:

```
1 >> fid = fopen (file_name, permission);
```

The `permission` string specifies the type of access you want to have:

- `'r'` - for reading only
- `'w'` - for writing only
- `'a'` - for appending only
- `'r+'` - both for reading and writing

Here is an example:

```
1 >> fid = fopen ('results.txt','w')    % tries to open the file results.txt
2                                     % for writing
```

The `fopen` statement returns an integer *file identifier*, which is a handle to the file (used later for addressing and accessing your file). When `fopen` fails (e.g. by trying to open a non-existing file), the file identifier becomes `-1`. It is also possible to get an error message, which is returned as the second optional output argument.

It is a good habit to test the file identifier each time when you open a file, especially for reading. Below, the example is given, when the user provides a string until it is a name of a readable file:

```

1 fid = 0;
2 while fid < 1
3     fname = input ('Open file: ', 's');
4     [fid, message] = fopen (fname, 'r');
5     if (fid == -1)
6         disp (message);
7     end
8 end

```

Exercise 20.1.

Create a script with the code given above and check its behavior when you give a name of a non-existing file (e.g. `noname.txt`) and a readable file (e.g. one of your functions). ■

When you finish working on a file, use `fclose` to close it up. MATLAB automatically closes all open files when you exit it. However, you should close your file when you finished using it:

```

1 fid = fopen ('results.txt', 'w');
2 ...
3 fclose(fid);

```

Type also `help fileformats` to find out which are readable file formats in MATLAB.

20.1 Text files

The `fprintf` command converts data to character strings and displays it on screen or writes it to a file. The general syntax is:

```

1 fprintf (fid,format,a,...)

```

For more detailed description, see section ???. Consider the following example:

```

1 >> x = 0:0.1:1;
2 >> y = [x; exp(x)];
3 >> fid = fopen ('exptab.txt', 'w');
4 >> fprintf(fid, 'Exponential function\n');
5 >> fprintf(fid, '%6.2f %12.8f\n', y);
6 >> fclose(fid);

```

Exercise 20.2.

Prepare a script that creates the `sintab.txt` file, containing a short table of the sinus function. ■

The `fscanf` command is used to read a formatted text file. The general function definition is:

```

1 [A,count] = fscanf (fid, format, size)

```

This function reads text from a file specified by file identifier `fid`, converts it according to the given `format` (the same rules apply as in case of the `fprintf` command) and returns it in a

matrix `A`. `count` is an optional output argument standing for the number of elements successfully read. The optional argument `size` says how many elements should be read from the file. If it is not given, then the entire file is considered. The following specifications can be used:

- `n` - read at most `n` elements into a column vector;
- `inf` - read at most to the end of the file;
- `[m,n]` - read at most `m*n` elements filling at least an `m`-by-`n` matrix, in column order; `n` can be `inf`.

Here is an example:

```
1 >> a = fscanf (fid, '%5d', 25);      % read 25 integers into a vector a
2 >> A = fscanf (fid, '%5d', [5 5]);    % read 25 integers into a 5 x 5 matrix A
```

MATLAB can also read lines from a formatted text and store it in a string. Two functions can be used for this purpose, `fgets` and `fgetl`. The only difference is that `fgetl` copies the newline character while `fgets` does not.

Exercise 20.3.

Create the following script and try to understand how it works (use the `help` command to learn more on the `feof` function):

```
1 fid = fopen ('exptab.txt', 'r');
2 title = fgetl (fid);
3 k = 0;
4 while feof(fid)~=1                % as long as end-of-file is not reached do
5     k = k+1;
6     line = fgetl (fid);            % get a line from the exponential table
7     tab(k,:) = str2num (line);     % convert the line into a vector from tab
8                                   % matrix
9 end
10 fclose(fid);
```

Look at the matrix `tab`. How does it differ from the originally created matrix? ■

Reading lines from a formatted text file may especially be useful when you want to modify an existing file. Then, you may read a specified number of lines of a file and add something at the found spot.

Exercise 20.4.

Create a script that reads the `exptab.txt` file and at the end of the file adds new exponential values, say, for $x = 1.1 : 0.1 : 3$. Note that you should open the `exptab.txt` file both for reading and writing. ■

A pair of useful commands to read and write ASCII delimited file (i.e. columns are separated by a specified delimiter such as space ' ' or tab, '\t') is `dlmread` and `dlmwrite`. A more general command is `textread`, which reads formatted data from a text file into a set of variables. Not only numeric data are read, but also characters and strings.

20.2 Working with Excel

Using `xlswrite` you write Microsoft Excel spreadsheet file (.xls)

This example writes the following mixed text and numeric data to the file `tempdata.xls`

```
1      d = {'Time', 'Temp'; 12 98; 13 99; 14 97};
```

The 4-by-2 matrix will be written to the rectangular region that starts at cell E1 in its upper left corner

```
1      s=xlswrite('tempdata.xls', d, 'Temperatures', 'E1')
```

Using `xlsread` you read Microsoft Excel spreadsheet file (.xls)

Exercise 20.5.

Prepare a script that creates the `costab.xls` file, containing a short table of the cosinus function.

■

Exercise 20.6.

Open the file `Mesh.inp`. Extract the coordinate of the nodes (they are introduced by the string `*Node`) and the nodal sets (they are introduced by the string `*Nset`).

Plot the nodes on a figure with circles and the nodes contained in the nodal sets with triangles.

■

Chapter 21


Numerical analysis

Numerical analysis can be used whenever it is impossible or very difficult to determine the analytical solution. This is often the case when analysing experimental results. MATLAB can be used to find, for example, the minimum, maximum or integral of a certain function or set of experimental data.

In this section we will discuss some issues on numerical integration and differentiation. Apart from that, the possibility to solve differential equations numerically will be introduced.

21.1 Differentiation

The derivative of a function $f(x)$ is defined as $f'(x) = \frac{df(x)}{dx} = \lim_{\Delta x \rightarrow 0} \frac{f(x+\Delta x) - f(x)}{(x+\Delta x) - x}$. In practice, the derivative of $f(x)$ at a is the slope of the line tangent to $f(x)$ at a . Numerical algorithms for computing the derivative of a function thus require the estimate of the slope of the function for some particular range of x values. Three common approaches to do this are the *backward difference*, *forward difference* and *central difference*. The figure below helps to illustrate the difference between these three approaches:



- **Backward difference:** $f'(x) \approx \frac{f(x_k) - f(x_{k-1})}{x_k - x_{k-1}}$
- **Forward difference:** $f'(x) \approx \frac{f(x_{k+1}) - f(x_k)}{x_{k+1} - x_k}$
- **Central difference:** $f'(x) \approx \frac{f(x_{k+1}) - f(x_{k-1}))}{x_{k+1} - x_{k-1}}$

To calculate the numerical derivative, the MATLAB function `diff(x)` can be used, which computes the differences between adjacent values of the vector `x`. If values of $f(x)$ are contained in vector `y` and the corresponding x values in vector `x`, the derivative of the function can be estimated using

```
1 >> deriv_y = diff(y)./diff(x);
```

The vector (or matrix) that `diff(x)` returns contains one less element (or row) than `x`. The corresponding x values are obtained from the original `x` vector by trimming either the first or last value;


- Trimming the last value results in a forward difference estimate.
- Trimming the first value results in a backward difference estimate.

To find and plot for instance the derivative of a function $f(x) = 5 \cos(10x) + x^3 - 2x^2 - 6x + 10$, we do:

```
1 >> x=0:.01:4;
2 >> y=5*cos(10*x) + x.^3 - 2*x.^2 - 6*x + 10;
3 >> subplot(1,2,1)
4 >> plot(x,y)
5 >> title('f(x) = 5cos(10x) + x^3 - 2x^2 - 6x +10')
6 >> deriv_y = diff(y)./diff(x);
7 >> xd = x(2:length(x)); % Backward difference x values
8 >> subplot(1,2,2)
9 >> plot(xd,deriv_y)
10 >> title('f`'(x)')
11 >> hold on
12 >> dy=-50*sin(10*x) + 3*x.^2 - 4*x - 6; % The exact function for f`'(x)
13 >> plot(xd,dy(2:length(x)),'r') % Showing that the numerical result is correct
```

21.2 Integration

The integral, or the surface underneath a 2D function, can be determined numerically. A broad family of algorithms for calculating the numerical value of a definite integral have been derived, some of which are displayed in the following figure.



- The left panel is a representation of the *Rectangle rule*. Here, the interpolating function is a constant function (a polynomial of degree zero) which passes through the point $((a+b)/2, f((a+b)/2))$. The approximation of the integral can be described as follows:

$$\int_a^b f(x)dx \approx (b-a)f\left(\frac{a+b}{2}\right)$$
- The middle panel is a representation of the *Trapezoidal rule*. Here, the interpolating function is a straight line (a polynomial of degree 1) which passes through the points $(a, f(a))$ and $(b, f(b))$. The approximation of the integral can be described as follows:

$$\int_a^b f(x)dx \approx (b-a)\frac{f(a)+f(b)}{2}$$
- The right panel is a representation of the *Simpson's rule*. Here, the interpolating function is a polynomial of degree 2 which passes through the end points a and b , and the midpoint $m = (a+b)/2$. The approximation of the integral can be described as follows:

$$\int_a^b f(x)dx \approx \frac{b-a}{6}(f(a) + 4f\left(\frac{a+b}{2}\right) + f(b))$$

As can be clearly seen in the figure, the accuracy of the different methods can be increased by dividing the interval (a, b) into multiple subintervals, computing an approximation for each subinterval, and then adding up all the results.

21.2.1 MATLAB commands

MATLAB has several commands to determine the integral of a function. One of these commands is `trapz`, which numerically integrates using the Trapezoidal rule. The accuracy of this method

also depends on the distance between the data points. Use `help` to learn more.

```

1 >> x = 0:0.5:10; y = 0.5 * sqrt(x) + x .* sin(x);
2 >> integrall = trapz(x,y)
3 integrall =
4     18.1655
5
6 >> x = 0:0.05:10; y = 0.5 * sqrt(x) + x .* sin(x);
7 >> integral2 = trapz(x,y)
8 integral2 =
9     18.3846

```

A more accurate result can be obtained by using the command `quad` or `quadl`, which numerically evaluate an integral of a function using an adapted Simpson's rule. Where the input for `trapz` is a vector or a matrix, the input for `quad` or `quadl` is specified by a **string** or by an **inline** definition (have a look at the intermezzo to learn more about inline functions).

Let $f = \frac{1}{(x-0.1)^2+0.1} + \frac{1}{(x-1)^2+0.1}$.

```

1 >> format compact
2 >> f='1./((x-0.1).^2+0.1)+1./((x-1).^2+0.1)';
3 >> integrall=quad(f,0,2)
4 integrall =
5     13.4118
6 >> integral2=quadl(f,0,2)
7 integral2 =
8     13.4118
9 >> whos
10  Name          Size          Bytes  Class      Attributes
11
12  f              1x37             74   char
13  integrall      1x1              8  double
14  integral2      1x1              8  double

```

Note in the previous example that `f` is defined as a character array (a string).

An extra parameter can be added to `quad` and `quadl` to specify the accuracy of the evaluation. If this parameter is not defined as in the previous example, the default error tolerance is $1.e-6$.

```

1 >> f='1./((x-0.1).^2+0.1)+1./((x-1).^2+0.1)';
2 >> integrall=quad(f,0,2)
3 integrall =
4     13.4118
5 >> integral2=quad(f,0,2,1.e-4)
6 integral2 =
7     13.4119
8 >> integral3=quad(f,0,2,1.e-1)
9 integral3 =
10     13.2654

```

Apart from the value of the approximated integral, MATLAB can also return the number of function evaluations that are needed to reach the defined (or default) error tolerance.

```

1 >> >> f='1./((x-0.1).^2+0.1)+1./((x-1).^2+0.1)';
2 >> [integrall,fcnt1]=quad(f,0,2)
3 integrall =

```

```

4      13.4118
5  fcnt1 =
6      97
7  >> [integral2,fcnt2]=quad(f,0,2,1.e-3)
8  integral2 =
9      13.4117
10 fcnt2 =
11      29

```

This clearly illustrates that larger values for the tolerance result in fewer function evaluations and thus faster computation.

INTERMEZZO

Inline functions

It may also be useful to define a function that will only be used during the current MATLAB session. MATLAB offers a command `inline` to define the so-called inline functions in the Command Window. For instance,

```

1 >> f = inline('cos(x).*sin(2*x)')
2 f =
3      Inline function:
4      f(x) = cos(x).*sin(2*x)
5 >> g = inline('sqrt(x.^2+y.^2)','x','y')
6 g =
7      Inline function:
8      g(x,y) = sqrt(x.^2+y.^2)

```

You can evaluate this function in a usual way:

```

1 >> f(-2)
2 ans =
3     -0.3149
4 >> g(3,4)
5 ans =
6      5
7 >> A = [1 2; 3 4];
8 >> B = [2 3; 4 5];
9 >> g(A,B)                                % the function also works with arrays
10 ans =
11      2.2361      3.6056
12      5.0000      6.4031

```

END INTERMEZZO

21.3 Solving differential equations - The ODE toolbox

Before going into mathematical details, we would like to point out that we have to be very careful using the ODE toolbox as a *black box solver* for various ordinary differential equations (ODEs). Note that you have to analyze carefully the *type* of ODE before you use a certain numerical method to solve it. With the numerical method, it is e.g. important to distinguish so-called *stiff* and *non-stiff* ODEs.

A *stiff* equation is a differential equation for which certain numerical methods for solving the equation are numerically unstable, unless the step size is taken to be extremely small.

Here, we discuss only functions which can be used for *non-stiff* problems, such as `ode23` and `ode45`. Note that in MATLAB also solvers for *stiff* problems are included; `ode15s`, `ode23s`, `ode23t` and `ode23tb`.

21.3.1 1st order ODE

MATLAB has an extensive library of functions for solving ordinary differential equations. In this course we will only show a very basic introduction. Therefore, we will focus on the main two, the built-in functions `ode23` and `ode45`, which implement versions of Runge-Kutta 2nd/3rd-order and Runge-Kutta 4th/5th-order, respectively.

Example 21.1:

Solve the following ODE of 1st order

$$\frac{dy}{dx} = x y^2 + y, \quad y_0 := y(x=0) = 1,$$

in the intervall $x \in [0, 0.5]$

For any differential equation in the form $dy/dx = y' = f(x, y)$, we begin by defining the function $f(x, y)$. Note that for single equations, we can define $f(x, y)$ as an inline function which simplifies the procedure. The basic usage for MATLAB's solver `ode45` is

```
1 ode45(function, domain, initial condition)
```

Thus, we can formulate the problem as

```
1 >> f = inline('x*y^2+y');
2 >> [x,y] = ode45(f,[0 0.5],1);
3 >> plot(x,y)
```

MATLAB returns two column vectors, the first with values of x and the second with values of y which can be plotted with the standard command `plot(x,y)`.

Approximating this solution numerically, the algorithm `ode45` has selected a certain discretization of the interval $[0, 0.5]$, and MATLAB has returned a value of y at each point in the partition. It is often the case that we would like to specify the partition, i.e. the numerical effort, on which MATLAB returns an approximation. For example, we might only want to approximate the solution at 5 points: $y(.1), y(.2), \dots, y(.5)$. We can specify this by entering the vector `[0, .1, .2, .3, .4, .5]` as the domain in `ode45`. That is, we use

```
1 >> ic = 1.0;
2 >> disc = 0:0.1:0.5;
3 >> f = inline('x*y^2+y');
4 >> [x,y] = ode45(f,disc,ic);
5 >> plot(x,y)
```

It is important to point out here that MATLAB continues to use roughly the same partition of values that it originally chose; the only thing that has changed is the values at which it is returning a solution. In this way, no accuracy is lost.

Several options are available for MATLAB's `ode45` solver, giving the user limited control over the algorithm. Two important options are relative and absolute tolerance, respectively `RelTol` and `AbsTol` in MATLAB. At each step of the `ode45` algorithm, an error is approximated for that step. If y_k is the approximation of $y(x_k)$ at step k , and e_k is the approximate error at this step, then MATLAB chooses its partition to insure

$$e_k \leq \max(\text{RelTol} * y_k, \text{AbsTol})$$

where the default values are `RelTol` = .001 and `AbsTol` = .000001. As an example for when we might want to change these values, observe that if y_k becomes large, then the error e_k will be allowed to grow quite large. In this case, we can increase the value of `RelTol`. For the equation $\frac{dy}{dx} = xy^2 + y$ with $y(0) = 1$, the values of y get quite large as x nears 1.

```
1 >> f=inline('x*y^2+y');
2 >> [x,y]=ode45(f,[0 1],1);
3 Warning: Failure at t=9.999897e-001. Unable to meet integration tolerances
4 without reducing the step size below the smallest value allowed (1.776357e-015)
5 at time t.
6 > In ode45 at 371
```

In order to fix this problem, we can choose a smaller value for `RelTol`.

```
1 >> f=inline('x*y^2+y');
2 >> options=odeset('RelTol',1e-10);
3 >> [x,y]=ode45(f,[0 1],1,options);
```

21.3.2 Systems of ODE

Solving a system of ODE in MATLAB is quite similar to solving a single equation, though since a system of equations cannot be defined as an inline function we must define it as an M-file. As an example, let's solve the system of Lorenz equations (The Lorenz equations have some properties of equations arising in atmospheric. Solutions of the Lorenz equations have long served as an example for chaotic behavior),

$$\frac{dx}{dt} = -\sigma x + \sigma y$$

$$\frac{dy}{dt} = \rho x - y - xz$$

$$\frac{dz}{dt} = -\beta z + xy$$

where for the purposes of this example, we will take $\sigma = 10$, $\beta = \frac{8}{3}$, and $\rho = 28$, as well as $x(0) = -8$, $y(0) = 8$, and $z(0) = 27$.

First, a MATLAB M-file `lorenz.m` is created that contains the Lorenz equations

```

1 function xyz = lorenz(t,x);
2 %LORENZ: Defines the Lorenz equations.
3 sig=10;
4 beta=8/3;
5 rho=28;
6 xyz=[-sig*x(1) + sig*x(2); rho*x(1) - x(2) - x(1)*x(3); -beta*x(3) + x(1)*x(2)];

```

Observe that x is stored as $x(1)$, y is stored as $x(2)$, and z is stored as $x(3)$. Additionally, xyz is a column vector, as is evident from the semicolon following the first appearance of $x(2)$.

Now, the solution of this system of ODE over the interval $t=[0,20]$ can be found:

```

1 >> x0=[-8 8 27];
2 >> tspan=[0,20];
3 >> [t,xyz]=ode45(@lorenz,tspan,x0);

```

The output for this last command consists of a column t of times followed by a matrix xyz with three columns, the first of which corresponds with values of x at the associated times, and similarly for the second and third columns for y and z .

Exercise 21.1.

Create a figure with 4 subplots. The first three should contain the components of the lorenz equation as introduced in this chapter as a function of t over the interval $t=[0,20]$. The fourth should display the 'Lorenz attractor', which is a 3D plot of the three coordinates. ■

21.4 Exercises

Exercise 21.2.

Determine the first and second derivative of the line through (x,y) with $x = 4:13$ and $y = [12.84 \ 12.00 \ 7.24 \ -0.96 \ -11.16 \ -20.96 \ -27.00 \ -24.96 \ -9.56 \ 25.44]$. Make a figure containing the plot of the initial data, the first derivative, and the second derivative in separate sub-plots. ■

Exercise 21.3.

Make a figure containing 3 sub-plots. Plot the function $f(x) = 2x\cos(10x) - 3x^2 + 6x + 10$ with $x = 0:.01:4$ and its derivative in separate sub-plots. Next, find the location of local minima and maxima and plot these, together with the function $f(x)$, in the third sub-plot (Note: a local minimum or maximum is defined by a zero crossing of the derivative. In other words, these are the locations where the values of the derivate go from a negative to a positive value or vice versa). ■

Exercise 21.4.

Write a script that repeats the example displayed in the section on noise. However, add a column of plots showing the filtered data, the first derivative of the filtered data, and the second derivative of the filtered data (thus acquiring a figure with 3x3 plots). Also plot the exact derivatives of the function in the graphs as a control (you can use the MATLAB symbolic toolbox to find the derivatives). ■

Exercise 21.5.

Write scripts that compute the integral of the function $f(x) = \frac{1}{1+x^2}$ using the Rectangle rule, the Trapezoidal rule and the Simpson's rule. Use the scripts to find the integral over the interval $[0, 4]$. Exercise with different accuracies (number of subintervals) and compare the results with the analytical value of the integral. ■

Exercise 21.6.

Adapt the scripts you prepared for the Trapezoidal rule and the Simpson's rule to determine the integral of $f(x) = \sin(x)$ over the interval $[0, \frac{\pi}{2}]$. Do this with different numbers of subintervals; $\mathbf{n}=[2 \ 4 \ 8 \ 16 \ 32 \ 64 \ 128 \ 256]$. For both the Trapezoidal rule and the Simpson's rule, create a matrix with 4 rows containing the values of n , the approximation of the integral, the error of the approximation, and the ratio between the errors of consecutive approximations. What does this tell you about the accuracy of both methods? ■

Exercise 21.7.

Find the integral of the function $f(x) = e^{-x^2/2}$ over the interval $[-3, 3]$. Exercise with the MATLAB commands `trapz`, `quad` and `quadl`, and different accuracies. With `quad` and `quadl`, get an idea of the number of function evaluations that are needed to reach a certain accuracy ■

Exercise 21.8.

The commands `quad` and `quadl` can also be used to integrate a function stored in a .m file (use `help` to find out how to do this). Make a .m file defining the function $y = \sin(x^5 + x^3)$. Then use `quad` to determine the integral of this function over the interval $[0, 2]$. Also try to find the exact value of the integral using the symbolic toolbox. Why can numerical integration be useful, even though it results in an approximation of the integral? ■

Exercise 21.9.

Solve the following ODE of 1st order

$$\frac{dy}{dx} = x y^2 + y, \quad y_0 := y(x=0) = 1,$$

in the domain $x \in [0, 0.9]$ with an equidistant discretization of $\Delta x = 0.01$. Use a m-file `firstode.m` to define the function. Plot the result. ■

Exercise 21.10.

The predator-prey model is a model that describes the number of predators and preys over time. When the prey is defined as x and the predator is defined as y , the number of predators and preys can be defined with the following system of ODE:

$$\frac{dx}{dt} = \alpha x(1 - y)$$

$$\frac{dy}{dt} = \beta y(x - 1)$$

where for the purposes of this example, we will take $\alpha = 5$ and $\beta = 1$, as well as $x(0) = 3$ and $y(0) = 1$. Solve the predator-prey model over the interval $t=[0, 10]$ and make a plot displaying the populations of preys and predators in a single graph.

■

Chapter 22

Exercises Day 1

22.1 Matrices and Vectors

Exercise 22.1.

Create a vector containing 5 elements such that its components are equally spaced when you take the logarithm of it ■

Exercise 22.2.

Perform the following exercises:

- Create a vector \mathbf{x} with the elements:
 - 1, 1/2, 1/3, 1/4, 1/5
 - 0, 1/2, 2/3, 3/4, 4/5

To do this, divide a vector \mathbf{y} by a vector \mathbf{z} .

- Given a vector \mathbf{t} , write down the MATLAB expressions that will compute:
 - $\ln(2 + \mathbf{t} + \mathbf{t}^2)$
 - $\cos(\mathbf{t})^2 - \sin(\mathbf{t})^2$
 - $e^{\mathbf{t}}(1 + \cos(3\mathbf{t}))$
 - $\tan^{-1}(\mathbf{t})$

Test them for $\mathbf{t} = 1 : 0.2 : 2$.

■

Exercise 22.3.

Use the knowledge on computing the inner product to find:

1. the Euclidean length of the vector $\mathbf{x} = [2 \ 1 \ 3 \ 7 \ 9 \ 4 \ 6]$, which is defined as $\|\mathbf{x}\| = \sqrt{(\Sigma \mathbf{x}_i^2)}$.

2. the angle between two column vectors, which is defined as $\cos \alpha = \frac{\mathbf{x}^T \mathbf{y}}{\|\mathbf{x}\| \|\mathbf{y}\|}$. Note that you can also use the MATLAB command `norm(v)`, which gives you the Euclidean length of a vector v . The MATLAB command `acosd(x)` can be used to find the inverse cosine of x expressed in degrees. Compute the angle between two vectors:

- $\mathbf{x} = [3 \ 2 \ 1]$ and $\mathbf{y} = [1 \ 2 \ 3]$
- $\mathbf{x} = 1 : 5$ and $\mathbf{y} = 6 : 10$

■

Exercise 22.4.

Clear all variables (use the command `clear`). Define the matrix $\mathbf{A} = [1:4; 5:8; 1 \ 1 \ 1 \ 1]$. Predict and check the result of the following operations:

- $\mathbf{x} = \mathbf{A}(:, 3)$
- $\mathbf{B} = \mathbf{A}(1 : 3, 2 : 2)$
- $\mathbf{A}(1, 1) = 9 + \mathbf{A}(2, 3)$
- $\mathbf{A}(2 : 3, 1 : 3) = [0 \ 0 \ 0; 0 \ 0 \ 0]$
- $\mathbf{A}(2 : 3, 1 : 2) = [1 \ 1; 3 \ 3]$
- $\mathbf{y} = \mathbf{A}(3 : 3, 1 : 4)$
- $\mathbf{A} = [\mathbf{A}; 2 \ 1 \ 7 \ 7; 7 \ 7 \ 4 \ 5]$
- $\mathbf{C} = \mathbf{A}([1, 3], 2)$
- $\mathbf{D} = \mathbf{A}([2, 3, 5], [1, 3, 4])$
- $\mathbf{D}(2, :) = []$

■

Exercise 22.5.

Define the matrices $\mathbf{T} = [3 \ 4; 1 \ 8; -4 \ 3]$ and $\mathbf{A} = [\text{diag}(-1:2:3) \ \mathbf{T}; -4 \ 4 \ 1 \ 2 \ 1]$. Perform the following operations on the matrix \mathbf{A} :

- extract a vector consisting of the 2nd and 4th elements of the 3rd row
- find the minimum of the 3rd column
- find the maximum of the 2nd row
- compute the sum of the 2nd column
- compute the mean of the 1st and 4th rows
- extract the submatrix consisting of the 1st and 3rd rows and all columns
- extract the submatrix consisting of the 1st and 2nd rows and the 3rd, 4th and 5th columns
- compute the total sum of the 1st and 2nd rows
- add 3 to all elements of the 2nd and 3rd columns

■

Exercise 22.6.

Let $\mathbf{A} = [2 \ 4 \ 1; 6 \ 7 \ 2; 3 \ 5 \ 9]$. Provide the commands which:

- assign the first row of \mathbf{A} to a vector \mathbf{x} ;
- assign the last 2 rows of \mathbf{A} to a vector \mathbf{y} ;
- add up the columns of \mathbf{A} ;
- add up the rows of \mathbf{A} ;



Exercise 22.7.

Let $A = [2 \ 7 \ 9 \ 7; \ 3 \ 1 \ 5 \ 6; \ 8 \ 1 \ 2 \ 5]$. Explain the results or perform the following commands:

- | | | |
|-----------------------------|---------------------------|---------------------------------------------------------------|
| • A' | • $\text{sum}(A)$ | • $\max(\min(A))$ |
| • $A(1,:)'$ | • $\text{sum}(A')$ | • $[[A; \text{sum}(A)] [\text{sum}(A, 2); \text{sum}(A(:))]]$ |
| • $A(:, [14])$ | • $\text{mean}(A)$ | • assign the even-numbered columns of A to an array B |
| • $A([23], [31])$ | • $\text{mean}(A')$ | • assign the odd-numbered rows to an array C |
| • $\text{reshape}(A, 2, 6)$ | • $\text{sum}(A, 2)$ | • convert A into a 4-by-3 array |
| • $A(:)$ | • $\text{mean}(A, 2)$ | • compute the reciprocal of each element of A |
| • $\text{flipud}(A)$ | • $\min(A)$ | • compute the square-root of each element of A |
| • $\text{fliplr}(A)$ | • $\max(A')$ | • remove the second column of A |
| • $[A \ A(\text{end}, :)]$ | • $\min(A(:, 4))$ | • add a row of all 1's at the beginning and at the end |
| • $[A; A(1:2, :)]$ | • $[\min(A)' \ \max(A)']$ | • swap the 2nd row and the last row |



Exercise 22.8.

Given the vectors $x = [1 \ 3 \ 7]$, $y = [2 \ 4 \ 2]$ and the matrices $A = [3 \ 1 \ 6; \ 5 \ 2 \ 7]$ and $B = [1 \ 4; \ 7 \ 8; \ 2 \ 2]$, determine which of the following statements can be correctly executed (and if not, try to understand why) and provide the result:

- | | | | |
|-------------------|-------------|--------------|---------------------------|
| • $x + y$ | • $[x; y']$ | • $B * A$ | • $B ./ x'$ |
| • $x + A$ | • $[x; y]$ | • $A . * B$ | • $B ./ [x' \ x']$ |
| • $x' + y$ | • $A - 3$ | • $A' . * B$ | • $2/A$ |
| • $A - [x' \ y']$ | • $A + B$ | • $2 * B$ | • $\text{ones}(1, 3) * A$ |
| • $[x; y] + A$ | • $B' + A$ | • $2 . * B$ | • $\text{ones}(1, 3) * B$ |



22.2 Visualization

Exercise 22.9.

Make a plot connecting the coordinates: (2, 6), (2.5, 18), (5, 17.5), (4.2, 12.5) and (2, 12) by a line. ■

Exercise 22.10.

Plot the function $y = \sin(x) + x - x \cos(x)$ in two separate figures for the intervals: $0 < x < 30$ and $-100 < x < 100$. Add a title and axes description. ■

Exercise 22.11.

Plot a circle with the radius $r = 2$, knowing that the parametric equation of a circle is

$[x(t), y(t)] = [r \cos(t), r \sin(t)]$ for $t = [0, 2\pi]$. ■

Exercise 22.12.

Plot an ellipse with semiaxes $a = 4$ and $b = 2$. ■

Exercise 22.13.

Plot the functions $f(x) = x$, $g(x) = x^3$, $h(x) = e^x$ and $z(x) = e^{x^2}$ over the interval $[0, 4]$ on the normal scale and on the log-log scale. Use an appropriate sampling to get smooth curves. Describe your plots by using the functions: `xlabel`, `ylabel`, `title` and `legend`. ■

Exercise 22.14.

Make a plot of the functions: $f(x) = \sin(1/x)$ and $f(x) = \cos(1/x)$ over the interval $[0.01, 0.1]$. How do you create x so that the plots look sufficiently smooth? ■

Exercise 22.15.

Produce a nice graph which demonstrates as clearly as possible the behavior of the function $f(x, y) = \frac{xy^2}{x^2+y^4}$ near the point $(0, 0)$. Note that the sampling around this points should be dense enough. ■

Exercise 22.16.

Plot a sphere, which is parametrically defined as $[x(t, s), y(t, s), z(t, s)] = [\cos(t) * \cos(s), \cos(t) * \sin(s), \sin(t)]$ for $t, s = [0, 2\pi]$ (use `surf`). Make first equal axes, then remove them. Use shading `interp` to remove black lines (use `shading faceted` to restore the original picture). ■

Exercise 22.17.

Plot the parametric function of r and θ : $[x(r, \theta), y(r, \theta), z(r, \theta)] = [r \cos(\theta), r \sin(\theta), \sin(6 \cos(r) - n\theta)]$ for $\theta = [0, 2\pi]$ and $r = [0, 4]$. Choose n to be constant. Observe, how the graph changes depending on different n . ■

Exercise 22.18.

Plot the surface $f(x, y) = xy e^{-x^2-y^2}$ over the domain $[-2, 2] \times [-2, 2]$. Find the values and the locations of the minima and maxima of this function. ■

Exercise 22.19.

Make a 3D smooth plot of the curve defined parametrically as: $[x(t), y(t), z(t)] = [\sin(t), \cos(t), \sin^2(t)]$ for $t = [0, 2\pi]$. Plot the curve in green, with the points marked by circles. Add a title, description of axes and the grid. You can rotate the image by clicking `Tools` at the Figure window and choosing the `Rotate 3D` option or by typing `rotate3D` at the prompt. Then by clicking at the image and dragging your mouse you can rotate the axes. Exercise with this option. ■

Exercise 22.20.

Plot an equilateral triangle with two vertices $[a \ a]$ and $[b \ a]$. Find the third vertex. Use `fill` to paint the triangle. ■

Chapter 23

Exercises Day 2

Exercise 23.1.

A simple pendulum is one which can be considered to be a point mass suspended from a string or rod of negligible mass. Let us consider oscillations with small amplitudes, so that $\sin \theta \simeq \theta$ and the position of the pendulum is described by

$$\theta = \theta_0 \cos \left(\sqrt{g/L} t \right).$$

To do:

1. Let the user enter the pendulum length and the initial position (angle θ) of the pendulum (Hint: look at the Matlab command `inputdlg`)
2. For these input values plot the variation of the angle θ as function of time
3. Plot the position of the pendulum at time $t = [T/4, T/2, 3T/4, T]$, where T denotes the period of the pendulum
4. Create a movie to illustrate the motion of the pendulum as shown in `MoviePendulum1.avi`.
5. Create a movie to illustrate the motion of the pendulum as shown in `MoviePendulum2.avi`.
Note that the color of the mass depends on its horizontal position.

■


Dont have the figures for this

Chapter 24


Exercises Day 3

24.1 Exercise 1


1. Open the file *Input3D.inp*. Extract the coordinate of the nodes and the element topology.
2. Plot the nodes on a figure with markers



3. Plot all the nodes that are on the external faces of the prismatic body




4. Plot the 8 vertices of the prismatic body and fill the top and bottom surfaces




5. Plot all the elements with filled tetrahedra

These are the Elements




6. Open the file *j_Input3D.dat*. Read the displacements of all the nodes and the stress for element 1 (For this particular type of loading the stress is the same in all the elements) at different time steps
7. Plot the stress for element 1 versus $\log((U_3 + 3)/3)$ (this is called true strain), where U_3 is the displacement in third direction of the node 141.




8. Plot the nodes at the last time step

These are the nodes at the last time increment



9. Plot the elements at the last time step

These are the Elements at the last time increment



10. Make a movie to capture the evolution of the mesh with time (see *Movie3D.avi*)
11. Make a movie to capture the evolution of U_3 (see *MovieU3.avi*). Assign to each element a displacement value \bar{U}_3 equal to the averaged value of the displacements U_3 of all its nodes.

24.2 Exercise 2

Write a code that allows the user to draw a geometrical figure on a figure panel using a series of points chosen by the user by clicking with the mouse.

- Open a figure and let the user select interactively the limits for the axes (use *inputdlg*). Adjust the limits for the axes accordingly.
- Let the user select interactively the colour of the lines for the drawing and the type of markers.
- the use should be able to draw on the figure panel using the mouse. By clicking the left mouse button he picks points. Draw a marker at the selected points and connect the points with a line. By clicking the right mouse button he picks the last point and stops drawing.
- Save all the points selected by the user in an excel file.

