

CS_191 Functional Programming I

Ulrich Berger

Department of Computer Science
Swansea University
Spring 2011

u.berger@swan.ac.uk , <http://www.cs.swan.ac.uk/~csulrich/>
tel 513380, fax 295708, room 306, Faraday Tower

1 Introduction

This course gives a first introduction to functional programming. No prior programming knowledge is assumed. The emphasis of the course is not on a full coverage of a particular programming language, but on general functional programming skills which can be applied in any functional and most main-stream programming languages.

The course is mainly based on material from the following books:

- Simon Thompson. *Haskell: The Craft of Functional Programming*, 2nd edition. Addison-Wesley, 1999.
- Graham Hutton. *Functional Programming in Haskell*. Cambridge University Press, 2007.
- Richard Bird. *Introduction to Functional Programming using Haskell*, 2nd edition. Prentice Hall, 1998.
- Paul Hudak. *The Haskell School of Expression – Learning Functional Programming through Multimedia*. Cambridge University Press, 2000.
- Bryan O’Sullivan, Don Stewart, and John Goerzen. *Real World Haskell*. O’Reilly, 2008. <http://book.realworldhaskell.org>
- Christoph Lüth. Praktische Informatik 3 (lecture slides). University of Bremen¹.

More information about books and online tutorials on functional programming can be found on the course web page <http://www.c.swa.ac.uk/~csulrich/fp1.html>

¹developed as part of the MMiSS project (MultiMedia Instruction in Safe Systems) and financially supported by the German Ministry for Education and Research

1.1 The fundamental ideas of functional programming

What is a functional program?

- A functional program is an *expression*.
Example: $(2 * 3) + 4$
- To *run* a functional program means to *evaluate* it.
Example: $(2 * 3) + 4$ evaluates to 10.
- There are *no assignments* (such as $x = x + 1$).

One can say that functional programming is similar to using a calculator.

Expressions are built from *functions* and other expressions by *function application*. In the example above, the expression $(2 * 3) + 4$ consists of the function addition $+$ applied to the expressions $2 * 3$ and 4. The expression $2 * 3$ in turn consists of the function multiplication $*$ applied to the expressions 2 and 3.

As with advanced calculators it possible to define new functions. For example, one can define the squaring function

```
square x = x * x
```

which can then be used, for example, in the expression `square (square (9 + 1))`.

Defining new functions from old ones is the main activity in functional programming, hence the name.

Here is an example highlighting the difference between functional and other styles of programming.

Example: Summing up the integers 1 to 10

In Java:

```
total = 0;
for (i = 1; i <= 10; ++i)
total = total+i;
```

The computation method is *assignment*.

In the functional language Haskell:

```
sum [1..10]
```

The computation method is *function application*

In the Haskell expression `sum [1..10]`

- `sum` is a function that computes the sum of the elements of a given list,
- the expression `[1..10]` evaluates to the list containing the integers 1 to 10,
- the whole expression `sum [1..10]` is the application of the function `sum` to the expression `[1..10]`, it evaluates to the integer 55.

Let us try to draw an analogy between different programming styles and building a house:

An imperative (C, Java, ...) program correspond to detailed instructions of what to *do* to build a house and in which *order*.

A functional program on the other hand corresponds to the architect's *construction plan*. Details of the execution of the program (plan) are left to the compiler (builder).

For example, in the functional expression $(3 * 4) + (2 * 5)$ it is neither specified which of the subexpressions, $3 * 4$ or $5 * 2$, is evaluated first, nor is it specified where the intermediate results are stored. These details are simply left to the compiler.

A very important property of functional programs is that, because there are no assignments or other side effects, *the values of variables are never changed*. In other words, *there is no state*. Therefore, *the value of a complex expression depends only on the values of its subexpressions* and nothing else. Due to this fact it is much easier to prove the *correctness* of a functional program than that of an imperative program.

Strengths of functional programming

- Simplicity and clarity (easy to learn, readable code)
- Reliability (correctness of a program can be proven),
- Productivity (algorithmic problems can be solved quickly).
- Maintainability and adaptability (programs can easily be adapted to changing requirements without introducing errors).
- Ericsson measured an improvement factor in productivity and reliability of between 9 and 25 in experiments on telephony software.

Because of these advantages functional programming is becoming more and more widespread. Typical application areas are Artificial Intelligence, Scientific computation, Theorem proving, Program verification, Safety critical systems, Web programming, Network toolkits and applications, XML parsing, Natural Language processing and speech recognition, Data bases, Telecommunication, Graphic programming, Finance and trading.

1.2 From Lambda-calculus to functional programming

In this course we use the functional programming language *Haskell*.

<http://www.haskell.org/>

Haskell is named after *Haskell B Curry*, an American Mathematician and major contributor to the theoretical foundation of functional programming.

Haskell B Curry (1900 – 1982)



Source: <http://www-groups.dcs.st-and.ac.uk/~history>

Curry and Alonzo Church (1903-1995) developed two closely related systems, *Combinatory Logic* and *Lambda-calculus*, which are based on the fundamental concept of a *function*. Their motivation was to build an alternative foundation for mathematics (which is usually founded on the notion of a *set*).

Today, Combinatory Logic and Lambda-calculus play an important role in many branches of theoretical computer science and they are the conceptual basis of functional programming.

We will not study Combinatory Logic and Lambda-calculus in this course, but if you have understood functional programming you have understood large parts of these theories.

Short history of functional programming

- Foundations 1920–30

Combinatory Logic and Lambda-calculus (Schönfinkel, Curry, Church)

- First functional languages 1960

LISP (McCarthy), ISWIM (Landin)

- Further functional languages 1970–80
 - FP (Backus); ML (Milner, Gordon), later SML and CAML; Hope (Burstall); Miranda (Turner)
- 1990: Haskell, Clean

Haskell has a polymorphic type system (variable types), a lazy semantics (this means that the evaluation of subexpressions is delayed as long as possible), and input/output is modelled by monads (in essence this means that expressions performing input/output are tagged and therefore clearly distinguished from other expressions).

The main current functional programming languages

Language	Types	Evaluation	I/O
<i>Lisp, Scheme</i>	type free	eager	via side effects
<i>ML, CAML</i>	polymorphic	eager	via side effects
<i>Haskell, Clean</i>	polymorphic	lazy	via monads

1.3 Organisation of the course

The lectures for this course will be supported by weekly lab classes where students may ask questions and test their skills by solving simple practical exercises. These exercises have to be completed and are assessed during that class. In the lab exercises and the courseworks it is recommended to *work in pairs*. This means i.p. that a coursework submission may be signed by up to two students. In this case it is assumed that both students have made approximately equal contributions.

Lab classes and exams

- **Lab classes**

Attendance is compulsory and will be monitored.

Lab exercises are compulsory and are assessed (10%).

Place and time: Will be announced in the lectures.

- **Coursework**

There will be **2 courseworks**, each worth 10%.

Deadlines are strict! Solutions will be handed out on the first Thursday after each submission date.

Submission: please put printout with signature(s) in the wooden box on the second floor. The box will be emptied in the morning after the submission date.

2 Getting started

This section explains how to work interactively with Haskell. Furthermore, some basic principles of Haskell's syntax are discussed.

2.1 Interactive Haskell

ghci

We will work with an interactive versions of Haskell. This makes it particularly easy to learn the language and test programs.

- ghci (Glasgow Haskell Compiler Interactive) comes with the GHC distribution. It is installed on the Linux computers in the labs.

Information on how to download, install and start ghci and can be found on the course web page.

We start ghci and evaluate some expressions:

```
>ghci
  _ _ _ _ _
 / _ \ / \ / \ / \ ( )
 / / \ \ / / \ / / | |      GHC Interactive, version 6.6.1, for Haskell 98.
 / / \ \ / \ / / \ \ | |      http://www.haskell.org/ghc/
 \_ _ _ / \ / \ \_ _ _ / | |      Type :? for help.
```

```
Loading package base ... linking ... done.
```

```
Prelude> 2 + 3
```

```
5
```

```
Prelude> (2 + 3) * 5
```

```
25
```

```
Prelude> max 4 7
```

```
7
```

2.2 Haskell scripts

We can write definitions of expressions into a *script*, that is, a text file and load the script.

The file must have the extension `.hs` (for “Haskell script”), for example `test.hs`

```
x :: Int
x = 2 + 3

y :: Int
y = x * 10
```

Definitions in Haskell

A Haskell definition consists of
a *type declaration* (or signature)

```
x :: Int
the defining equation
x = 2 + 3
```

Both, type declaration and defining equation must start at the beginning of a line!

Loading and testing a script.

```
Prelude> :load test.hs
[1 of 1] Compiling Main
Ok, modules loaded: Main.
Main> x
5
Main> y
50
Main> x - y
-45
Main> x * x
25
Main> z
<interactive>:1:0: Not in scope: 'z'
```

Developing a script and terminating a session.

We can extend our script by typing into the file `test.hs` more definitions:

```
area :: Float
area = 1.5^2 * pi
```

Before testing the new definitions we have to reload the file

```
Main> :reload
Main> area
7.068583
Main>:quit
```

- The command `:quit` terminates the session.
- The command `:?` lists all available commands.
- Commands can be abbreviated by their first letters.

A good script contains plenty of comments.

```
-- My first Haskell script
-- Here are my first Haskell definitions

a :: Float
a = 3.5

b :: Float
b = -7.1

{- Here I try then out (copy and paste from the terminal)
*Main> a
3.5
*Main> b
-7.1
*Main> a + b
-3.6
*Main> b^2
-}
-- End of script
```

2.3 Functions

Defining a function

We can define a *function* that computes the area of a circle for a given radius `r`.

```
ar :: Float -> Float
ar r = r^2 * pi
```

The signature `ar :: Float -> Float` means that `ar` is a function that expects a floating point number as input and returns a floating point number as output.

The variable `r` is called *formal parameter*. It represents an arbitrary input.

Instead of "input of a function" one often says "argument of a function".

```
*Main> ar 1.5
7.0685835
```

Functions of more than one argument

```
average :: Float -> Float -> Float
average x y = (x+y)/2
```

The signature `average :: Float -> Float -> Float` means that `average` is a function that expects two floating point numbers as input (represented by the formal parameters `x` and `y`) and returns a floating point number as output.

```
*Main> average 3 4
3.5
*Main> average (3,4)
```

```
<interactive>:1:8:
  Couldn't match expected type 'Float' against inferred type ...
```

The second function call is incorrect since the function `average` expects two arguments which must be separated by spaces (no brackets, no comma).

Exercise 1 Define a function that computes the average of three numbers.

Exercise 2 Define a function that computes the hypotenuse c of a right-angled triangle from its catheti (legs) a and b using the formula

$$c = \sqrt{a^2 + b^2}$$

You may use the predefined Haskell function `sqrt :: Float -> Float` that computes the square root of a non-negative floating point number.

Exercise 3 Define a function `tri-area` that computes the area of the triangle above from its catheti a, b .




Figure 1: A right-angled triangle with catheti a, b and hypotenuse c .

Exercise 4 Define a function `tri-circum` that computes the circumference of the triangle above from its catheti a, b .

Exercise 5 Define a function `rect-area` that computes the area of a rectangle above from its sides a, b .

Exercise 6 Define a function `rect-circum` that computes the circumference of a rectangle above from its sides a, b .

Exercise 7 Visit the course web page

<http://www.cs.swan.ac.uk/~csulrich/fp1.html>

in particular check out the links

Instructions on how to get started in the labs

Haskell predefined functions and type classes

2.4 Syntactic issues

While in mathematics the arguments of a function are usually enclosed in brackets and separated by commas, in Haskell arguments are written without brackets and separated by spaces. Brackets are only used to resolve syntactic ambiguities, for example to distinguish the expression $(3+4)*5$ from $3 + (4 * 5)$.

Mathematical notation versus Haskell syntax

Maths	Haskell
$\max(x, y)$	<code>max x y</code>
$3(x^2 + y^2) + 5$	<code>3 * (x^2 + y^2) + 5</code>
$\text{ar} : \text{Float} \rightarrow \text{Float}$	<code>ar :: Float -> Float</code>
$\text{ar}(a, b) = \sqrt{a^2 + b^2}$	<code>ar a b = sqrt (a^2 + b^2)</code>

The preference rules for mathematical operations are as usual (multiplication binds stronger than addition, exponentiation binds stronger than multiplication). For example, the expression $(3^2) + (4 * 5)$ can be written equivalently as $3^2 + 4 * 5$.

Haskell ignores spaces between operators with special names (such as `+`, `*`, `^`) and arguments. For example the expressions `3 + 4` and `3+4` are the same for Haskell.

Case sensitivity, prefix vs. infix

- Haskell is case sensitive:

identifiers, function names and formal parameters must be lower cases (`x` and `area`, but not `X` and not `Area`)

types must be upper case (`Float`, but not `float`)

- Prefix operators bind stronger than infix operators:

`max 3 4 + 5` means `(max 3 4) + 5`

but not `max 3 (4 + 5)`.

Which editor should I use for writing a Haskell script?

In principle any text editor can be used. It is important that a font is used which displays each character (including space) with equal width.

A highly recommended editor is **Emacs** (or XEmacs). Emacs is installed on the Linux machines in the lab.

The use of emacs in the labs is compulsory.

Other editors, in particular inadequate ones like Notepad will not be tolerated.

Instructions on how to use Emacs and how to download Emacs on your own computer can be found on the course web page.

Emacs recognises Haskell scripts and assists, for example, in checking brackets.

It is recommended not to use tabs.

Exercise 8 Familiarise yourself with the use of emacs in the labs.

Further Exercises

Exercise 9 Define a function `mark` that computes the final mark for `cs_191` from its components.

The function should take as as inputs

- The mark for the lab exercises (counts 10%),
- The mark for Coursework 1 (counts 10%),
- The mark for Coursework 2 (counts 10%),
- The mark for the exam (counts 70%).

All marks (for the components and the final mark) should be given in percent, that is, as a floating point number between 0 and 100.

Don't forget to define the signature of your function.

Exercise 10 Continuing Exercise 9, define a function `exam` that takes as inputs marks for the labs and the courseworks and computes the minimum mark sufficient for getting a final mark of at least 40%.

3 Types

In this Chapter we introduce Haskell's most common basic types and their associated predefined functions.

What are types?

A type is a name for a collection of similar objects.

There are

- *predefined primitive types* such as Bool, Int, Integer, Float, Double, Char (these will be discussed next)
- *complex types* such as pairs, lists, arrays, function types
- *user defined types*

Haskell is a *strongly typed* programming language, that is, Haskell offers type safety.

Haskell does *static type checking*, that is, types are checked during compile time.

For first users Haskell's type checker may appear a bit fussy (many type errors), but the strict typing discipline has many advantages.

Why types?

- Early detection of errors at compile time
- Compiler can use type information to improve efficiency
- Type signatures facilitate program development
- and make programs more readable
- Types increase productivity and security

3.1 Booleans

The type `Bool` of Boolean values (named after George Boole (1815-1864) an English mathematician and philosopher) contains exactly two values, `True` and `False`. It is the second simplest Haskell type. There is one even simpler type, `()`, called "unit type", which contains only one value. This type will be discussed later.

Booleans

- Values: `True` and `False` (uppercase!)
- Predefined functions:

<code>not</code>	<code>:: Bool -> Bool</code>	<i>negation</i> (not)
<code>(&&)</code>	<code>:: Bool -> Bool -> Bool</code>	<i>conjunction</i> (and)
<code>()</code>	<code>:: Bool -> Bool -> Bool</code>	<i>disjunction</i> (or)

Here is how these functions behave (“ \Rightarrow ” means “evaluates to”):

```

not True      => False
not False    => True

True && True   => True
True && False  => False
False && True  => False
False && False => False

True || True  => True
True || False => True
False || True => True
False || False => False

```

Example: Exclusive Or

The predefined function `||` defines *Inclusive Or*:

`x || y` evaluates to `True` if `x`, or `y`, or both are `True`.

We can define *Exclusive Or* using the predefined functions `||`, `&&` and `not`:

```

(|*) :: Bool -> Bool -> Bool
x |* y = (x || y) && not (x && y)

```

For example:

```

True |* False => True
True |* True  => False

```

but

```

True || True => True

```

Using infix operators in prefix notation and vice versa.

We may use the operator `|*|` as well in prefix notation. For this to work, the operator has to be enclosed in round brackets. For example, we may write `(|*|) True True` which for Haskell is the same as `True |*| True`.

On the other hand, the maximum operator `max` may be used in infix notation if enclosed by inverted commas, for example, `3 'max' 4` is equivalent to `max 3 4`.

Here is the general rule for infix and prefix notation:

Infix vs. prefix

A function of two arguments whose name is composed from the special symbols

`! , # , $, % , & , * , + , - , / , < , = , > , ? , @ , ^ , |`

for example `||`, can be used

- infix without brackets: `x || y`
- prefix with brackets: `(||) x y`

A function of two arguments whose name is composed of letters, for example `max`, can be used

- infix with inverted commas: `x 'max' y`
- prefix without inverted commas: `max x y`

Beware of the difference: `x 'max' y` (correct) vs. `x 'max' y` (incorrect).

In a type signature the prefix version of a name must be used:

```
(||) :: Bool -> Bool -> Bool.
```

Exercise 11 Define *Exclusive Or* using the name `xor` and using prefix notation throughout the definition.

Solution.

```
xor :: Bool -> Bool -> Bool
xor x y = (&&) ((||) x y) (not ((&&) x y))
```

Exercise 12 Define *Inclusive Or* using `not` and `&&`. That is, the defined function should have the same behaviour as the function `||`.

Solution.

```
(|||) :: Bool -> Bool -> Bool
x ||| y = not ((not x) && (not y))
```

Exercise 13 Define `&&` using `not` and `||`.

Give two definitions, one using infix notation, the other using prefix notation.

Note: you cannot use the names `&&` or `and` since these are names of (different) predefined functions.

Boolean functions are also called *logic gates*.

The boolean values `True` and `False` are called *constructors*. They can be used on the left hand side of the defining equations of a logic gate.

Defining a logic gate by its truth table

We can define the function `xor` by listing its truth table:

```
xor :: Bool -> Bool -> Bool
xor True  True   = False
xor True  False  = True
xor False True   = True
xor False False  = False
```

One may also use wildcards:

```
xor :: Bool -> Bool -> Bool
xor True  False = True
xor False True  = True
xor _     _     = False
```

If wildcards are used, then the *order* of the defining equations matters!

Exercise 14 Use wildcards to give short definitions of the logic gates `&&` and `||`.

Exercise 15 Define the logic gate `nand` which returns `False` exactly when both inputs are `True`. This gate is also known as the *Sheffer stroke* and is often denoted with a vertical bar.

Give two definitions of `nand`: from predefined gates and by listing its truth table.

Exercise 16 Define `&&` using `nand` only. You are not allowed to use the constructors `True` and `False` or any other logic gate.

Hint: Define `not` first.

3.2 Numbers

Basic numeric types

Computing with numbers

Limited precision constant cost	\longleftrightarrow	arbitrary precision increasing cost
------------------------------------	-----------------------	--

Haskell offers:

- `Int` - integers as machine words
- `Integer` - arbitrarily large integers
- `Rational` - arbitrarily precise rational numbers
- `Float` - floating point numbers
- `Double` - double precision floating point numbers

Some predefined numeric functions

```
(+), (*), (^), (-), min, max :: Int -> Int -> Int
-   :: Int -> Int           -- unary minus
abs :: Int -> Int           -- absolute value
div :: Int -> Int -> Int    -- integer division
mod :: Int -> Int -> Int    -- remainder of int. div.
gcd :: Int -> Int -> Int    -- greatest common divisor
```

```

3 ^ 4      ⇒ 81
-9 + 4     ⇒ -5      div 33 12 ⇒ 2
2-(9 + 4) ⇒ -11     mod 33 12 ⇒ 9
abs -3     ⇒ error   gcd 33 12 ⇒ 3
abs (-3)  ⇒ 3

```

Exercise 17 (a) Let X be a set containing m elements and let Y be a set containing n elements.

How many functions from X to Y do exist?

For example, if $X = Y = \text{Bool}$, then there are 4 such functions: `id` (identity, predefined, reproduces the input), `not` (negation, predefined, swaps True and False), `constTrue` (always returns True), `constFalse` (always returns False).

(b) How many n -ary logic gates do exist?

An n -ary logic gate is a function $f :: \underbrace{\text{Bool} \rightarrow \dots \rightarrow \text{Bool}}_n \rightarrow \text{Bool}$.

(c) Write a Haskell program that computes for a given input n the number of n -ary logic gates.

Solution.

(a) To construct an arbitrary function $f :: X \rightarrow Y$ one selects for every $x \in X$ one of the n elements of Y . Hence, there are $\underbrace{n * \dots * n}_m = n^m$ such functions

(b) Mathematically, an n -ary logic gate is the same as function from the set

$$\text{Bool}^n := \underbrace{\text{Bool} \times \dots \times \text{Bool}}_n$$

to the set `Bool`. Since Bool^n has 2^n elements and `Bool` has 2 elements, there are, according to part (a), $2^{(2^n)}$ such functions, that is, there are $2^{(2^n)}$ n -ary logic gates.

(c)

```

gates :: Int -> Int
gates n = 2^(2^n)

```

```

{-
*Main> gates 1
4
*Main> gates 2
16

```

```
*Main> gates 3
256
*Main> gates 4
65536
*Main> gates 5
0
-}
```

What went wrong?

Answer: $\text{gates } 5 = 2^{(2^5)} = 2^{32} = 65536^2$ which is too big for the type `Int`. The problem can be solved by replacing `Int` with `Integer`:

```
gates1 :: Integer -> Integer
gates1 n = 2^(2^n)

{-
*Main> gates1 4
65536
*Main> gates1 5
4294967296
*Main> gates1 10
1797693134862315907729305190789024733617976978942306572734300811577326
7580550096313270847732240753602112011387987139335765878976881441662249
2847430639474124377767893424865485276302219601246094119453082952085005
7688381506823424628814739131105408272371633505106845862982399472459384
79716304835356329624224137216
-}
```

Comparison operators

`(==), (/=), (<=), (<), (>=), (>) :: Int -> Int -> Bool`

But also

`(==), (/=), (<=), (<), (>=), (>) :: a -> a -> Bool`

where `a = Bool, Integer, Rational, Float, Double`

```
-9 == 4  => False
 9 == 9  => True
 4 /= 9  => True
 9 >= 9  => True
 9 > 9   => False
```

Floating point numbers: Float, Double

- Single and double precision Floating point numbers
- The arithmetic operations (+), (-), (*), / may also be used for Float and Double
- Float and Double support the same operations

```
(/) :: Float -> Float -> Float
pi  :: Float
exp, log, sqrt, logBase, sin, cos :: Float -> Float
```

```
3.4/2      ⇒ 1.7
pi         ⇒ 3.14159265358979
exp 1     ⇒ 2.71828182845905
log (exp 1) ⇒ 1.0
logBase 2 1024 ⇒ 10.0
cos pi    ⇒ -1.0
```

Exercise 18 The formula for the roots of a quadratic is

$$\frac{-b \pm \sqrt{b^2 - 4ac}}{2a}$$

Define functions `smallerRoot`, `largerRoot` which return the smaller and larger root of a quadratic. The coefficients a, b, c shall be given as floats and the result shall be a float. You may assume that the quadratic does have roots ($b^2 - 4ac \geq 0$) and is not degenerate ($a \neq 0$).

One can switch to double precision floating point numbers by using the type `Double`.

```
Prelude> pi :: Float
3.1415927
Prelude> pi :: Double
3.141592653589793
```

Note that the constant `pi` is overloaded. It can be used as a value of type `Float` and of type `Double`. The overloading is resolved by explicit type annotation.

All constants and functions for floating point numbers are overloaded in this way.

Exercise 19 Define two versions of the function `sqrt`. One should operate on `Float`, the other on `Double`.

Solution.

```
sqrt1 :: Float -> Float
sqrt1 x = sqrt x

sqrt2 :: Double -> Double
sqrt2 x = sqrt x
```

```
{-
*Main> sqrt1 2
1.4142135
*Main> sqrt2 2
1.4142135623730951
-}
```

Conversion from and to integers

```
fromIntegral :: Int -> Float
fromIntegral :: Integer -> Float

round :: Float -> Int    -- round to nearest integer
round :: Float -> Integer
```

Example

The following code does not compile:

```
half :: Int -> Float
half x = x / 2
```

The reason is that division (`/`) expects two floating point numbers as arguments, but `x` has type `Int`.

Correct code:

```
half :: Int -> Float
half x = (fromIntegral x) / 2
```

Exercise 20 (a) Define a function that computes the average of three integers, rounded *down* to the nearest integer.

(b) Define a function that computes the average of three integers, rounded to the nearest integer.

Solution.

```
avround :: Int -> Int -> Int -> Int
avround x y z = round (fromIntegral (x+y+z) / 3)
```

```
avdown :: Int -> Int -> Int -> Int
avdown x y z = (x+y+z) `div` 3
```

```
{-
*Main> avround 3 4 4
4
*Main> avdown 3 4 4
3
-}
```

3.3 Characters and strings

Notation for characters: 'a'

Notation for strings: "Hello"

```
(:) :: Char -> String -> String -- prefixing
(++) :: String -> String -> String -- concatenation
```

(:) binds stronger than (++)

```
"Hello " ++ 'W' : "orld!" ⇒ "Hello World!"
```

Is the same as

```
"Hello " ++ ('W' : "orld!")
```

Exercise 21 (a) Define a function `rep` that repeats a string (that is, concatenates it with itself). Evaluate the expressions `rep "hello! "`, `rep (rep "hello! ")`, and so on.

- (b) Define a function `rep2` that applies the function `rep` to a string twice, that is, `rep2 x = rep (rep x)`. How many repetitions of the string "hello! " do we get when evaluating the expression `rep2 (rep2 (rep2 "hello! "))` ?

The ASCII code of a character

```
fromEnum :: Char -> Int
```

computes the ASCII code of a character:

```
fromEnum 'a' ⇒ 97
fromEnum 'b' ⇒ 98
fromEnum '9' ⇒ 57
```

```
toEnum :: Int -> Char
```

is its inverse. `toEnum 97 :: Char ⇒ 'a'`

The ASCII codes of characters denoting capital letters, small letters and digits are arranged in blocks. For example, the characters 'A', ..., 'Z' have the codes 65, ..., 90.

Exercise 22 Define a function that transforms a capital letter into a small letter.

Solution. Since the codes of letters are arranged in blocks we can simply compute the code of a capital letter shift the code by a certain offset and compute the corresponding character.

```
offset :: Int
offset = fromEnum 'a' - fromEnum 'A'

toLower :: Char -> Char
toLower x = toEnum (fromEnum x + offset)
```

Exercise 23 Define a function that tests whether a character is a digit, that is, one of the characters '0', ..., '9'.

Exercise 24 Define a function that tests whether an integer is the ASCII code of a lower case vowel. The lower case vowels are the letters a, e, i, o, u.

4 Control and structure

In this chapter we learn how to

- control the execution of a program by various forms of *definition by cases* (if-then-else, guarded equations, case-expressions, pattern matching),
- handle *exceptions*,
- structure programs by *local definitions* (`let`, `where`),
- organise large systems of programs into *modules* and use Haskell's *libraries*,
- make programs more *readable* by introducing, *type synonyms*.

4.1 Definition by cases

If-then-else

```
max2 :: Int -> Int -> Int
max2 x y = if x < y then y else x
```

```
signum :: Int -> Int
signum x = if x > 0 then 1 else
           if x == 0 then 0 else (-1)
```

The general rule for forming an if-then-else expression is:

If e is an expression of type `Bool` and
 e_1, e_2 are expressions of the same type a ,
then `if e then e1 else e2` is an expression of type a .

Note that if-then-else builds an ordinary expression which can be used at any place where the type fits. For example

```
(if 3 < 2 then 4 else 5) * 6
```

is a well-formed expression. What is its value?

Exercise 25 Define a function that computes the absolute value of a floating point number (note that this function is predefined and has name `abs`).

Exercise 26 Define a function that computes the maximum of three integers.

The nesting of the “if-then-else” in the definition of `signum` can be expressed more elegantly using *guarded equations*.

Guarded equations

```
signum1 :: Int -> Int
signum1 x
  | x > 0      = 1
  | x == 0     = 0
  | otherwise  = 1
```

Note that `otherwise` is a constant that has value `True`.

Exercise 27 Define a function that takes as inputs three integers m , n , x , and computes one of the strings "far left", "left", "middle", "right", or "far right" depending on whether

$$x < m,$$

$$x = m,$$

$$m < x < n,$$

$$x = n, \text{ or}$$

$$x > n.$$

Case analysis with pattern matching

```
empty :: String -> Bool
empty s = case s of
    ""    -> True
    _     -> False
```

```
tl :: String -> String
tl xs = case xs of
    ""      -> ""
    (_:xs) -> xs
```

Layout. The clauses in a case-expression (the lines with the arrow in the middle) must be *left aligned*. For example, this code does *not* compile

```
tl :: String -> String
tl xs = case xs of
    ""      -> ""
    (_:xs) -> xs
```

but that one does

```
tl :: String -> String
tl xs = case xs of
    ""      -> ""
    (_:xs) -> xs
```

The Haskell layout rules can be switched off by enclosing the clauses in curly brackets and separating them by semicolons.

```
empty :: String -> Bool
empty s = case s of {""    -> True ; _ -> False}

tl :: String -> String
tl xs = case xs of {"" -> "" ; (_:xs) -> xs}
```

Booleans or Strings as pattern

Case-expressions can be also used with Boolean constants or constant strings as pattern:

```
(&&&) :: Bool -> Bool -> Bool
x &&& y = case x of {False -> False ; True -> y}

age :: String -> Int
age x = case x of
    {"John" -> 21 ; "Alice" -> 20 ; _ -> -1}
```

The functions (`&&&`) and `age` can be equivalently defined by a list of defining equations.

Defining equations with pattern matching

```
(&&&) :: Bool -> Bool -> Bool
False &&& y = False
_      &&& y = y

age :: String -> Int
age "John"    = 21
age "Alice"  = 20
age _        = -1
```

Exercise 28 Define a function `swap12` that swaps the first two characters of a string provided the string has at least length 2. Otherwise reproduce the string unchanged. Define `swap12` in two ways: using a case-expression and using a list of defining equations.

Solution.

```
swap12 :: String -> String
swap12 (x:y:s) = y:x:s
swap12 t      = t

swap12 :: String -> String
swap12 t = case t of {(x:y:s) -> y:x:s ; _ -> t}
```

The general rule for forming case-expressions is as follows: If

- e is an expression type a ,
- p_1, \dots, p_n are patterns of type a ,
- e_1, \dots, e_n are expressions of type b ,

then `case e of { $p_1 \rightarrow e_1$; ... ; $p_n \rightarrow e_n$ }` is an expression of type b .

A *pattern* is an expression built from constructors and different variables.

Constructors are constants such as `True`, `False`, `0`, `1`, ..., `"`, and `:` (attaching a character in front of a string. Note that a string like `"hi"` is syntactic sugar for the pattern `'h':'i':"`.

A case-expression is evaluated as follows:

1. the expression e is evaluated,
2. the first pattern p_i matching this value is selected,
3. the expression e_i is evaluated and returned as the value of the case-expression,
4. if no match is possible an error is raised.

Note that a pattern may contain wildcards (`_`) since different occurrences of the wildcard are just syntactic sugar for different variables.

Exercise 29 Define a function that contracts the first two characters of a string into one provided these two characters are equal and the string has at least length 2. Otherwise reproduce the string unchanged.

4.2 Exception handling

A better solution for the function `age` is to raise an exception when the person is not known:

```
age :: String -> Int
age "John"    = 21
age "Alice"   = 20
age x = error ("The age of " ++ x ++ " is not known")
```

The general formation rule for error-expressions is:

- If e is an expression of type `String`,
- then `error s` is an expression of type a for any type a .
- When an error-expression is evaluated, then the computation of the whole expression where it occurs is aborted and the string s is displayed.

Exercise 30 Modify the function in Exercise 28 such that if the string does not have at least length 2, then an exception with an appropriate error message is raised.

4.3 Local definitions

`let and where`

```
g :: Float -> Float -> Float
g x y = (x^2 + y^2) / (x^2 + y^2 + 1)
```

better

```
g x y = let s = x^2 + y^2
         in s / (s + 1)
```

or

```
g x y = s / (s + 1) where
  s = x^2 + y^2
```

Exercise 31 Recall that in Exercise 18 we computed the roots x of a quadratic equation

$$ax^2 + bx + c = 0$$

from the coefficients a, b, c according to the formula

$$x_{1,2} = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}$$

Improve the solution by structuring the code and raising exceptions when no root exists or infinitely many root exist. The function that computes the roots should have an extra Boolean argument that allows one to choose between the two roots.

Solution.

```

roots :: Float -> Float -> Float -> Bool -> Float
roots a b c p

  | a /= 0      = let d = b^2 - 4*a*c
                  in if d >= 0
                      then ((-b) + k * sqrt d) / (2*a)
                      else err

  | a == 0      = if b /= 0 then (-c)/b else err

where k  = if p then 1 else (-1)
      err = error "no root or infinitely many roots"

```

In order to test whether our function works correctly we define a function that computes $ax^2 + bx + c$ for given a, b, c, x .

```

quad :: Float -> Float -> Float -> Float -> Float
quad a b c x = a*x^2 + b*x + c

```

```

{-
*Main> roots 1 2 (-3) True
1.0
*Main> quad 1 2 (-3) 1
0.0
-}

```

In order to make testing more convenient we define a test function.

```

testroots :: Float -> Float -> Float -> Bool
testroots a b c =
  let {x1 = roots a b c True ; x2 = roots a b c False}
      in quad a b c x1 == 0 && quad a b c x2 == 0

```

```

{-
*Main> testroots 1 2 (-3)
True
*Main> testroots 1 5 (-3)
False
-}

```

What's wrong?

How can we fix the error?

Exercise 32 Write a program that tests the correctness of the function `sqrt`.

Make sure that the test cannot be applied to negative numbers by an appropriate exception handling. Soften your test so that the expected results are obtained. How “hard” can you make your test?

Local definition of a function

The sum of the areas of two circles with radii `r`, `s`.

```
totalArea :: Float -> Float -> Float
totalArea r s = pi * r^2 + pi * s^2
```

We make the program more modular and readable by using an auxiliary function to compute the area of one circle:

```
totalArea :: Float -> Float -> Float
totalArea r s = circleArea r + circleArea s

where circleArea x = pi * x^2
```

Note that `circleArea` is a function from `Float` to `Float`.

We could add the signature of `circleArea` to make this fact explicit:

```
totalArea :: Float -> Float -> Float
totalArea r s = circleArea r + circleArea s

where circleArea :: Float -> Float
      circleArea x = pi * x^2
```

Exercise 33 Define the function `totalArea` using `let`.

Exercise 34 Define a variation of the solution of Exercise 31 where the line

```
then ((-b) + k * sqrt d) / (2*a)
```

is replaced by

```
then ((-b) +- sqrt d) / (2*a)
```

and the function (+-) is defined locally.

Solution.

```
roots :: Float -> Float -> Float -> Bool -> Float
roots a b c p
```

```
  | a /= 0    = let d = b^2 - 4*a*c
                 in if d >= 0
                    then ((-b) +- sqrt d) / (2*a)
                    else err
```

```
  | a == 0    = if b /= 0 then (-c)/b else err
```

where

```
(+-) = if p then (+) else (-)
```

```
err  = error "no root or infinitely many roots"
```

Exercise 35 [Suggested by a student]

Write a program computing the *standard deviation* of five values.

Note that the standard deviation of x_1, \dots, x_5 is given by the formula

$$\sqrt{\frac{(x_1 - m)^2 + (x_2 - m)^2 + (x_3 - m)^2 + (x_4 - m)^2 + (x_5 - m)^2}{5}}$$

where m is the mean value of x_1, \dots, x_5 .

Use local definitions of functions and values to obtain a short and readable program.

Remark. Later we will learn how to compute the standard deviation of an arbitrary *list* of values.

Solution.

```
stdev :: Float -> Float -> Float -> Float -> Float -> Float
stdev x1 x2 x3 x4 x5 = sqrt (mean (d x1) (d x2) (d x3) (d x4) (d x5))
```

where

```
mean y1 y2 y3 y4 y5 = (y1 + y2 + y3 + y4 + y5)/5
```

```
m = mean x1 x2 x3 x4 x5
```

```
d x = (x-m)^2
```

4.4 Modules

Modules for large program systems

Local definitions are useful for structuring small program units.

In order to structure large program systems one uses *modules*.

- A module must have a name beginning with a capital letter, for example
Geometry
- The module must reside in a file with the same name, but with extension `.hs`.
In our example
Geometry.hs

A module can look, for example, as follows:

```
module Geometry

-- Exported functions
(
```

```
rectangleArea,      -- :: Float -> Float -> Float
rectangleCircumference, -- :: Float -> Float -> Float
circleArea,        -- :: Float -> Float
circleCircumference -- :: Float -> Float
)
```

where

```
-- Definitions of exported functions
```

```
rectangleArea :: Float -> Float -> Float
rectangleArea a b = a*b
```

```
rectangleCircumference :: Float -> Float -> Float
rectangleCircumference a b = 2*(a+b)
```

```
circleArea :: Float -> Float
circleArea r = pi*r^2
```

```
circleCircumference :: Float -> Float
circleCircumference r = 2*pi*r
```

- If the list of exported functions is omitted, then all functions defined in the module are exported.
- In order to use a module in a Haskell script one has to import it.

For example

```
import Geometry
```

```
-- imported functions
(
  rectangleArea,
  circleArea
)
```

```
twocircleArea :: Float -> Float -> Float
twocircleArea r1 r2 = circleArea r1 + circleArea r2
```

- If the list of imported functions is omitted, then all functions defined in the module are imported.

Exercise 36

Create a module that contains test functions for Haskell’s implementation of the square root function (`sqrt`) and the natural logarithm (`log`). Note that the following equations are expected to hold.

$$\begin{aligned}(\sqrt{x})^2 &= x && \text{for } x \geq 0 \\ e^{\log(x)} &= x && \text{for } x > 0\end{aligned}$$

The Haskell function mapping x to e^x is called `exp`. Your test functions should have an extra parameter controlling the “hardness” of the test.

Import your module in another Haskell script and try it out.

On the Haskell web page one finds a large collection of library modules. They can be used in the same way as our example module. For example, by writing

```
import Char
```

we import the Haskell library `Char` with all its functions. These functions are concerned with operations on characters.

It is also possible to import a module directly from the command line.

```
Prelude> isControl '\n'
<interactive>:1:0: Not in scope: ‘isControl’
Prelude> :module Char
Prelude Char> isControl '\n'
True
Prelude Char> isControl 'a'
False
Prelude Char> :t isControl
isControl :: Char -> Bool
```

4.5 Type synonyms

Programs become more readable if types are given names that indicate what its members are used for.

For example, in the definitions of the areas of a rectangle and a circle the formal parameters of type `Float` sometimes play the role of the side of a rectangle and sometimes the role of the radius of a circle. The results represent areas.

```
type Side    = Float
type Radius  = Float
type Area    = Float

rectangleArea :: Side -> Side -> Area
rectangleArea a b = a*b

circleArea :: Radius -> Area
circleArea r = pi*r^2
```

Remark. With the keyword `type` one does not define a new type. One just introduces a new name for an existing type.

Exercise 37

Introduce type synonyms to make the program computing the roots of a quadratic more readable.

5 Input/Output

So far we studied “pure” functional programs which can be evaluated, but which don’t interact with the “outside world”.

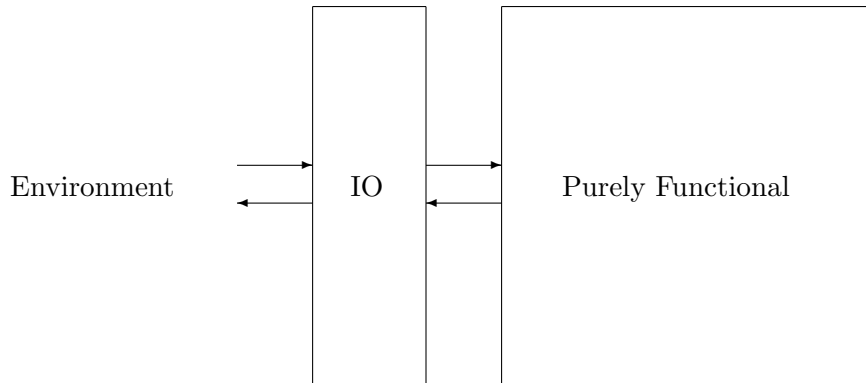
Of course, *some* interaction *does* happen when we run a program:

- (1) the expression we type is read from the terminal as a string;
- (2) the string is parsed, evaluated and the value transformed back into a string;
- (3) that string is printed at the terminal, and haskell is ready to read the next expression.

This so-called *read-eval-print loop* is carried out by the interactive Haskell program `ghci`. The actions (1) and (3) are typical examples of *Input/Output*, *IO*, for short.

We now discuss how to write simple IO-programs and how to combine them with purely functional programs in a sensible way.

IO-programs aren’t purely functional, but they are necessary to run and use purely functional programs. It is the philosophy of functional programming to strictly separate the purely functional and the IO part of a program. The functional part should do the main computation (and it should be the by far bigger part), whereas the IO-part acts as an interface between the functional part and the environment. The advantage of this separation is that the largest and complicated part of a program can be made highly reliable because it can, for example, be *proven* to be correct. On the other hand, the IO-part should be kept as small and simple as possible because it is there where errors (in the hardware or software, or by the user) are most likely to occur and hardest to detect and eliminate. The separation of the functional and the IO part can be done, in principle, in any programming language, but functional programming languages particularly encourage and support it.

IO-program as interface between environment and purely functional program**5.1 Writing to and reading from the terminal**

We begin with simple interaction with the terminal.

Printing a string at the terminal

```
hw :: IO ()
hw = putStrLn "Hello, world!"
```

- The type `IO ()` is the type of *actions*.
- `putStrLn :: String -> IO ()` is a pre-defined Haskell function that takes a string as input and performs the action of printing it at the standard output (usually terminal or console) and inserting a new line.
- Evaluating `hw` will print `Hello, world!` on the terminal.

Carrying out several actions in sequence

```
helloGoodbye :: IO ()
helloGoodbye = do putStrLn "Hello world!"
                  putStrLn "Goodbye world!"
```

Equivalent notation with curly brackets and semi-colons:

```
helloGoodbye :: IO ()
helloGoodbye =
  do {
    putStrLn "Hello world!" ;
    putStrLn "Goodbye world!"
  }
```

For `do`-expressions the same layout rules apply as for `case`- and `let`-expressions. If the curly-bracket-semi-colon syntax is used no layout rules apply. (Recall that this is the same as for `case`- and `let`-expressions.)

Reading a string from the terminal

```
echo :: IO ()
echo = do {
  putStr "Please enter a string: " ;
  s <- getLine ;
  putStrLn s ;
  putStrLn s
}
```

- `getLine :: IO String` reads a line from the standard input (usually the console or terminal) and returns it as a string. A line is usually entered by the user and terminated by pressing the Return-key.
- `IO String` is the type of actions that return a string as result.
- The line `s <- getLine` means that the line read from the terminal is bound to the local variable `s` which then can be used in subsequent clauses of the `do`-expression.

Remarks.

1. The variable `s` is local to the `do`-expression and therefore not visible outside of it.
2. The line `s <- getLine` is similar to an assignment in imperative languages. However, the difference is that `s` is a *new* local variable. When executing `s <- getLine` it never happens that an existing variable is overwritten.
3. In the program `echo` the two clauses printing `s` at the terminal could be contracted to one that prints the concatenation of `s` with itself plus the newline character `'\n'` in the middle:

```
echo1 :: IO ()
echo1 = do {
    putStr "Please enter a string: " ;
    s <- getLine                    ;
    putStrLn (s ++ "\n" ++ s)
}
```

4. One can use `getLine` repeatedly:

```
echo2 :: IO ()
echo2 = do {
    putStr "Please enter a string: " ;
    s <- getLine                    ;
    putStrLn s                      ;
    putStr "and another one, please: " ;
    t <- getLine                    ;
    putStrLn (s ++ " " ++ t)
}
```

5.2 The type `IO a`

- For any type `a`, the type `IO a` is the type of actions (side-effects) that in addition return a value of type `a`.
- The only way to access the value returned by an action `e :: IO a` is by a clause `x <- e` within a `do`-expression.
- It is *impossible* to define a function `getVal :: IO a -> a` that would “extract” the value returned by an action.
- As a consequence, an expression that has a type not containing any `IO` is purely functional and does not perform any action (that is, side effect). In particular, its value does not depend on the environment.
- Therefore, the type `IO a` visibly separates the action part of a program from the purely functional part.

Remark. In most (virtually all) other programming languages `IO`-types do not exist. This means, in effect that, for example, the types `IO String` and `String` are identified. If this were the case in Haskell, then, we would have, for example,

```
getLine :: String
```

and we could perform the test

```
getLine == getLine
```

The result would not always be the expected value `True`, but depend on the evaluation strategy and on the strings entered by the user.

This shows that without the distinction between the types `String` and `IO String` even the most basic reasoning about programs (for example the law that everything is equal to itself) breaks down.

The type `IO ()`

- `()` is the type containing exactly one element which is again denoted `()`.
- Hence, the type `IO ()` is the type of actions that do not return any useful value.
- Typically, the main program is of type `IO ()` because its only purpose is to do some action.
- The type `()` is often called “unit type”. It corresponds to Java’s type `void`.

Exercise 38 Program an interactive English-French dictionary.

First Solution.

```
dictionary1 :: IO ()
dictionary1 =
  do en <- getLine
     case en of
     {
       "apple"  -> putStrLn "pomme"  ;
       "bread"  -> putStrLn "pain"   ;
       "butter" -> putStrLn "beurre" ;
       "milk"   -> putStrLn "lait"   ;
       -       -> putStrLn "sorry, don't know"
     }
}
```

This works, but it is bad programming style since actions (`putStrLn ... :: IO ()`) are scattered all over the program.

Second Solution.

```

dictionary2 :: IO ()
dictionary2 = do {
    en <- getLine ;
    putStrLn (en2fr en)
}

en2fr :: String -> String
en2fr en =
  case en of
  {
    "apple"  -> "pomme"   ;
    "bread"  -> "pain"    ;
    "butter" -> "beurre"  ;
    "milk"   -> "lait"    ;
    _       -> "sorry, don't know"
  }

```

This solution is better. It clearly divides the program in a small interactive part (`dictionary2`) and a large purely functional part (`en2fr`).

Note that in order to see that the program `en2fr` is purely functional it suffices to look at its type (`String -> String`) which has no `IO` in it. It is not necessary to check the code (which might be large, in practice).

Exercise 39 Write the program `en2fr` using

- (a) pattern matching with several equations,
- (b) `if-then-else`,
- (c) guarded equations.

The program `dictionary2` can be made slightly more readable using a `let`-clause:

```

dictionary3 :: IO ()
dictionary3 = do {
    en <- getLine      ;
    let {fr = en2fr en} ;
    putStrLn fr
}

```

Note that the `let`-clause has no “`in`”.

5.3 Using recursion to create an interaction loop

We can modify our dictionary so that one can enter as many queries as one likes until the empty string is entered:

```
dictloop :: IO ()
dictloop = do {
    en <- getLine ;
    if en == "" then return ()
      else do {
                putStrLn (en2fr en) ;
                dictloop
            }
}
```

- `return ()` is the action that does nothing. Here it is used to terminate the interaction.
- In general, the function `return :: a -> IO a`, where `a` is any type, transforms any value `x` of type `a` into an action that does nothing but returning that value `x`.
- The looping behaviour is achieved by a *recursive call* of `dictloop`.
- In Haskell, *recursion* is the only means for creating loops.

Exercise 40 Write a program that collects nonempty strings entered at the terminal and prints them all if the user enters the empty string.

In an imperative language we would solve this by writing a loop with a variable that holds the words entered so far and which is updated by the new word entered in each round. At the end the content of the local variable would be returned as result.

Here is how we do this in a functional language?

```
collect :: IO ()
collect = do {
    allwords <- coll "" ;
    putStrLn allwords
}
where
```

```
coll :: String -> IO String
coll history = do {
    s <- getLine          ;
    case s of
        "" -> return history
        _  -> coll (history ++ " " ++ s)
}
```

We see that the local variable is modelled by the local *function* `coll` which is recursively called with an “updated” argument.

```
*Main> collect
when
is
this
lecture
over?

    when is this lecture over?
*Main>
```

5.4 Input/Output with numbers

Suppose we want to compute the solutions of a quadratic whose coefficients are entered through the terminal. The solutions shall be displayed at the terminal.

Because at the terminal we can only enter strings, we need a function that transforms strings into numbers:

```
read :: String -> Float
```

Furthermore, in order to display the computed roots on the terminal, we need a function that translates numbers into strings:

```
show :: Float -> String
```

Recall that we did use the function `show` earlier.

The functions `read` and `show` can be used for other number types and types such as `Bool` or `Char` as well.

Solving a quadratic interactively

```

solve :: IO ()
solve = do {
    sa <- getLine ;
    sb <- getLine ;
    sc <- getLine ;
    let {
        a  = read sa ;
        b  = read sb ;
        c  = read sc ;
        x1 = roots a b c True ;
        x2 = roots a b c False ;
        sc1 = show x1 ;
        sc2 = show x2
    } ;
    putStrLn ("The solutions are " ++ sc1 ++ " and " ++ sc2)
}

```

```

roots :: Float -> Float -> Float -> Bool -> Float
roots a b c p

```

```

| a /= 0    = let { d = b^2 - 4*a*c }
               in if d >= 0
                   then ((-b) + k * sqrt d) / (2*a)
                   else err

```

```

| a == 0    = if b /= 0 then (-c)/b else err

```

```

where k      = if p then 1 else (-1)
      err = error "no root or infinitely many roots"

```

```

*Main> solve
2
5
-7
The solutions are 1.0 and -3.5
*Main> solve
-5
8
20
The solutions are -1.3540659 and 2.9540658

```

Note that the computationally interesting part is completely covered by the purely functional program `roots`. The main program `solve` only handles the interaction with the terminal.

Exercise 41 Write a program that computes the length of a string entered by the user.

You may use the predefined function

```
length :: String -> Int
```

Exercise 42 Modify the program of exercise 41 such that the user can compute lengths of strings repeatedly (see program `echoLoop`).

Exercise 43 Write a program that computes the sum of numbers entered by the user (see program `collect`).

Exercise 44 Modify the program of exercise 43 such that the user can compute sums repeatedly.

Exercise 45 Modify the program `solve` such that the user is asked whether the solutions should be checked and, if that's the case, the user may specify the required accuracy of the solutions.

5.5 Communication with files

Reading from and writing to files can be done with the operations

- Reading a file:

```
readFile    :: FilePath -> IO String
```

- Writing into a file (override, append):

```
writeFile   :: FilePath -> String -> IO ()  
appendFile :: FilePath -> String -> IO ()
```

where `FilePath` is a predefined type synonym:

```
type FilePath = String
```

Example: Displaying the length of a file

```
lengthFile :: IO ()
lengthFile =
  do {
    fn <- getLine ;
    file <- readFile fn ;
    putStrLn ("File " ++ fn ++ " contains " ++
              show(length file) ++ " characters.\n")
  }
```

```
*Main> lengthFile
fp1-io.tex
File fp1-io.tex contains 23723 characters.
```

Example: Reversing the content of a file and writing it into another one

```
reverseFile :: FilePath -> FilePath -> IO ()
reverseFile fn1 fn2 =
  do {
    str <- readFile fn1 ;
    writeFile fn2 (reverse str)
  }
```

Note that the above program expects the filenames as arguments. Hence a valid call of `reverseFile` is, for example

```
reverseFile "28-2.hs" "test.hs"
```

Notice that the filenames have to be enclosed in double quotes because we use `ghci`'s predefined `IO`.

Exercise 46 Write a program that writes for a given integer n all numbers $1 \dots n$ into a specified file.

Exercise 47 Write a program that checks whether two files have the same content.

5.6 Random data

The Haskell module `Random`, which is loaded by including the line

```
import Random
```

at the beginning of a Haskell script, has several operations for generating random numbers. For example,

```
randomRIO :: (Random a) => (a,a) -> IO a
```

The type of `randomRIO` says that for any type in the type class `Random` (for example, `Int`, `Char`) and any two values `x`, `y` of that type the expression

```
randomRIO (x,y)
```

evaluates to an action returning an element in the range between `x` and `y`.

Example: Generating a random password.

```
password :: IO ()
password = do {
    let {c = randomRIO ('a','z')} ;
        c1 <- c ;
        c2 <- c ;
        c3 <- c ;
        c4 <- c ;
        putStrLn [c1,c2,c3,c4]
    }
```

Exercise 48 Improve the password program such that the length of the password can be determined by the user.

Exercise 49 Using the solution to Exercise 48, write a program that generates random passwords of random lengths between 1 and 10.

Exercise 50 Write a program that allows the user to enter names (as a strings) and that assigns to each name a random password (say of length 8) and writes names and passwords into a file.

Haskell has many more libraries (that is, modules) for other kinds of IO operations, for example, GUIs, communication with the operating system, and interfaces with other programming languages. Click the link *Haskell libraries* for more information.