

Concepts of Programming Languages

Lecture Notes: Functional Programming

Stefan Mitsch

School of Computing, DePaul University
smitsch@depaul.edu

Scala is a multi-paradigm language that combines functional and object-oriented programming. In this lecture, we discuss pattern matching, list processing functions, and compare the difference between methods and functions in Scala in more depth.

PATTERN MATCHING

When expressing contracts, we often want to describe properties about the resulting data. These can be as simple as comparing values (e.g., $x \geq 0$), but usually require us to describe the expected shape of the result in more detail.

For example, when faced with the following description of how to process a list, a large language model may generate code that we want to safeguard with a contract.

>_ In Scala, implement a method that takes a list of at least three numbers and returns it with the first three sorted in ascending order.

 Large language model response:

```
1 def f (numbers: List[Int]) : List[Int] = {
2   require(numbers.length >= 3, "The list must contain at least 3
      elements")
3   val (firstThree, rest) = numbers.splitAt(3)
4   val sortedFirstThree = firstThree.sorted
5   sortedFirstThree ++ rest
6 } ensuring { case xs => ??? }
```

The ??? operator throws a `NotImplementedException` and it is our task to describe our requirements more formally with a suitable contract. In this example, we are interested in testing the shape of the list and that the first three elements are in ascending order.

In a traditional imperative programming language style we often access the elements of the list by index like below, or using projections:

```
1 ensuring { case xs =>
2   val x1 = xs(0) // xs.head
3   val x2 = xs(1) // xs.tail.head
4   val x3 = xs(2) // xs.tail.tail.head
5   x1 <= x2 && x2 <= x3
6 }
```

This style results in lengthy contracts that do not clearly express the intended shape of the result. *Pattern matching* is an alternative way of decomposing a complex data structure into its elements: it emphasizes the shape of the result and simultaneously branches on shapes and binds variables.

```

1 ensuring {
2   case xs =>
3     val x1 :: x2 :: x3 :: Nil = xs
4     x1 <= x2 && x2 <= x3
5 }

```

Even more concisely, we can express the pattern immediately in the signature of the function that we pass to `ensuring` as follows:

```

1 ensuring {
2   case x1 :: x2 :: x3 :: Nil =>
3     x1 <= x2 && x2 <= x3
4 }

```

Pattern matching *binds variables* to the components of a data structure according to the specified shape. For example, when computing the sum of two numbers in a tuple of type `(Int, Int)`, a traditional approach may use a conditional expression to check for the tuple being non-null, and then accesses the tuple elements by their index as below:

```

1 def sum(p: (Int, Int)) : Int =
2   if p==null then throw MatchError(p)
3   val x = p(0)
4   val y = p(1)
5   x + y

```

Using pattern matching, we can express the same behavior by listing *cases* for the expected shapes, and in each case, listing the variables that stand for the components of the complex data structure.

```

1 def sum(p: (Int, Int)) = p match
2   case (x, y) => x+y

```

Types are optional:

```

1 def sum(p: (Int, Int)) = p match
2   case (x: Int, y: Int) => x+y

```

Pattern matching *branches* between multiple cases. This is useful when a data structure can take one of multiple shapes (e.g., an empty list vs. a list with at least one element) and we want to produce different results for those shapes. In the following example, we print the first element of a non-empty list; otherwise we print that the list is empty. In an imperative programming approach, we typically use a conditional expression to distinguish the cases. Each branch of the conditional expression then performs a different computation, as below:

```

1 def printHead(xs: List[Int]) : String =
2   if xs == Nil then "List is empty"
3   else
4     val y: Int = xs.head
5     val ys = xs.tail
6     s"List is non-empty, head is $y"

```

The same functionality can be achieved with pattern matching:

```

1 def printHead(xs: List[Int]) : String = xs match
2   case Nil           => "List is empty"
3   case (y: Int) :: ys => s"List is non-empty, head is $y"

```

The code matches on the shape of list `xs`. If `xs` is the empty list (`case Nil`), then we return the string `"List is empty"`; else (`case (y: Int) :: ys`), we have a list with head `y` and tail `ys` and we print the head. Unused variables (tail `ys` above) can be omitted from the pattern shape using the wildcard operator `_` as below:

```

1 def printHead(xs: List[Int]) = xs match
2   case Nil           => "List is empty"
3   case y :: _        => s"List is non-empty, head is $y"

```

Patterns can be nested to take any arbitrarily complicated shape:

```

1 def f (xs: List[(Int, String)]) = xs match
2   case Nil           => "List is empty"
3   case _ :: Nil      => "List has one element"
4   case _ :: (x, _) :: _ => s"The second int is $x"

```

The code above lists three cases:

- `case Nil` matches the empty list to return string `"List is empty"`
- `case _ :: Nil` matches a list with exactly one element
- `case _ :: (x, _) :: _` matches any list of at least two elements, and it binds the `Int` of the second element to variable `x`

FUNCTIONS OVER LISTS

We start by inspecting an example to print the elements of a list.

```

1 def printList (xs: List[Int]) : Unit = xs match
2   case Nil     => ()
3   case y :: ys =>
4     println(y) // can format: println("0x%02x".format(y))
5     printList (ys)
6
7 val xs = List(11, 21, 31)
8 printList (xs)
9 // 11 21 31

```

In the example above, we see two ways of printing the elements of the list (unformatted vs. formatted).

What if we now want to apply some other function to every element of the list? The basic setup of the recursive algorithm wouldn't change, only the specific operation that we apply at each element does. We can describe that abstract idea of processing every element with a recursive algorithm that, in addition to the list being processed, takes a function to be applied to each element as an argument.

```

1 def foreach (xs: List[Int], f: Int=>Unit) : Unit = xs match
2   case Nil     => ()
3   case y :: ys =>

```

```

4      f (y)
5      foreach (ys, f)
6
7  val xs = List(11, 21, 31)
8  foreach (xs, println)

```

Now it becomes easy to make variations.

```

1  def printHex (x: Int) = println("0x%02x".format(x))
2  foreach (xs, printHex)

```

But do we really care about the elements in the list? An additional improvement makes the element type a type parameter of the method.

```

1  def foreach [X] (xs: List[X], f: X=>Unit) : Unit = xs match
2      case Nil => ()
3      case y::ys =>
4          f (y)
5          foreach (ys, f)

```

Finally, we may not even always want to define the functions that we apply to each element. For this, Scala supports Lambda expressions (anonymous functions):

```

1  foreach (xs, (x: Int) => println("0x%02x".format(x)))
2  foreach (xs, println("0x%02x".format(_)))
3
4  // also possible
5  val printHex = (x: Int) => println("0x%02x".format(x))
6  foreach (xs, printHex)

```

We do not need to implement `foreach` ourselves, Scala collections provide it!

The examples above use both type and value parameters: type parameters are in square brackets, whereas value parameters are in round brackets. All type parameters must be declared before value parameters. Functions themselves are of function type: for example, `X=>Int` is the type of a function taking an argument of type `X` and returns a result of type `Int`. In Lambda expression, types are often unnecessary if Scala can infer them (type inference is smarter on methods than functions).

List comprehensions. From mathematics, we might be familiar with *set comprehensions* of the form

$$\{(m, n) \mid m \in \{0, \dots, 10\} \wedge n \in \{0, \dots, 10\} \wedge m \leq n\} .$$

List comprehensions of a similar form are included in many programming languages, such as SETL, Haskell, Scala, and JavaScript.

Scala provides another builtin special syntax to express `foreach` using list comprehensions (named *for-expressions* in Scala).

```

1  for x <- xs do println("0x%02x".format(x))

```

We are now inspecting a (less-than-optimal) way of expressing imperative loops with our `foreach` implementation, by using a variable in scope:

```

1 def sum (xs : List[Int]) : Int =
2     var result = 0
3     xs.foreach ((x : Int) => result = result + x)
4     result

```

Later, we'll see how *folds* provide a better way of expressing such ideas.

A note on equality: Java uses builtin operators for reference equality, and a method for value equality; Scala has methods for both, the operator symbol method `==` for value equality and method `eq` for reference equality.

Transformers. A frequent operation on collections is the modification of elements in the collection. To this end, *transformers* are functions to build a list of modified elements while traversing a collection recursively (unlike above where we only print elements but do not manipulate them).

```

1 def transform (xs : List[Int]) : List[String] = xs match
2     case Nil => Nil
3     case y::ys => ("0x%02x".format(y)) :: transform(ys)

```

A transformer is expected to take one cons cell as input and produce another cons cell as output. Just like `foreach`, there is a builtin way of applying transformers to collections: `map`.

Scala again provides special notation to apply transformers in a for-expression:

```

1 for x <- xs do println("0x%02x".format(x))
2 // is compiled to xs.foreach (x => println("0x%02x".format(x)))
3
4 for x <- xs yield "0x%02x".format(x)
5 // is compiled to xs.map (x => "0x%02x".format(x))

```

Filtering. Often, we want to apply a function only to elements satisfying a certain condition, omitting the remaining elements in the output collection.

```

1 def filter [X] (xs : List[X], f : X=>Boolean) : List[X] = xs match
2     case Nil => Nil
3     case y::ys if f(y) => y :: filter(ys, f)
4     case _::ys => filter(ys, f)
5
6 val zs = (0 to 7).toList
7 filter(zs, ((_:Int) % 3 != 0))

```

Again in special for-expression notation:

```

1 for z <- zs; if z % 3 != 0 yield z
2 // compiles to zs.filter (z => z % 3 != 0)
3
4 for z <- zs; if z % 3 != 0 yield "0x%02x".format(z)
5 // compiles to zs.filter (z => z % 3 != 0).map (z => "0x%02x".format(z))
6
7 for z <- zs; if z % 3 != 0 do println("0x%02x".format(z))
8 // compiles to zs.filter (z => z % 3 != 0).foreach (z => println("0x%02x".format(z)))

```

```

1 def flatten [X] (xs: List [ List [X] ]) : List [X] = xs match
2   case Nil    => Nil
3   case y::ys => y ::: flatten (ys)
4
5 val xss = List ( List (11, 21, 31), List (), List (41, 51) )

```

In the above implementation of `flatten`, we use operator `:::`.

Multiple iterators. For-expressions are quite general and can combine nested iterators in a single pass.

```

1 val xss = List ( List (11, 21, 31), List (), List (41, 51) )
2 for xs <- xss;
3   x <- xs yield (x, xs.length)
4 // result: List [(Int, Int)] = List ((11, 3), (21, 3), (31, 3),
5   (41, 2), (51, 2))

```

We can also compute the cross product of independent iterators.

```

1 val xs = List (11, 21, 31)
2 val ys = List ("a", "b")
3 for x <- xs;
4   y <- ys yield (x, y)
5 // result: List [(Int, String)] = List ((11, a), (11, b), (21, a),
6   (21, b), (31, a), (31, b))

```

CURRYING

We say that functions are *first-class* if they can be

- declared within any scope
- passed as arguments to other functions, and
- returned as results of functions.

Higher-order functions are functions that can take other functions as arguments (e.g., `map` etc.). Methods that take multiple arguments can be defined in the “usual” way with multiple formal arguments, or in a *curried* way as higher-order definitions that take one argument at a time and produce methods that take the remaining arguments.

```

1 // paired method
2 def add1(x: Int, y: Int) = x+y
3 add1(11, 21)
4
5 // curried method
6 def add2(x: Int)(y: Int) = x+y // takes an Int, returns method of
   type (y: Int): Int
7 add2(11)(21)
8
9 // paired function

```

```

10 val add3 = (x:Int , y:Int) => x+y
11
12 // curried function
13 val add4 = (x:Int) => (y:Int) => x+y // takes an Int, returns a
   function of type Int=>Int
14
15 // can mix notations: method returns a function
16 def add5(x:Int) = (y:Int) => x+y

```

Both paired and curried notation support partial function application:

```

1 // paired
2 val add1p = add1(4, _)
3 add1p(1) // result 5
4
5 // curried
6 val add4p = add4(4)
7 add4p(1) // result 5
8 add4p(2) // result 6

```

Functions and methods in Scala can also be created explicitly by instantiating the appropriate classes:

```

1 def a (x:Int) = x + 1;
2 val b = (x:Int) => x + 1;
3 val c = new Function[Int,Int] { def apply(x:Int) = x + 1 }
4 val d : PartialFunction[Any, Int] = { case i: Int => i + 1 }
5
6 val fs = List(a,b,c,d)
7 for f <- fs yield f(4)

```

In summary:

- `def` defines a method, parameter types explicit
- `=>` defines a function, parameter types inferable
- Functions are objects with method `apply(e( args ) ==>e.apply( args ))`

FOLDS

MapReduce is a programming model for processing and generating data sets with a parallel, distributed algorithm. It requires a main process that performs filtering and sorting (the map step) and a summary operation that collects and combines results (the reduce step). Below is an example of counting the number of occurrences of each word in a set of documents using MapReduce in Scala.

```

1 def map(name: String , contents: String) =
2   // name: document name (irrelevant here)
3   // contents: document content
4   for w <- contents do
5     emit (w, 1)
6

```

```

7 def reduce(word: String, partialCounts: Iterator) =
8   var sum = 0
9   for pc <- partialCounts do
10     sum = sum + pc
11   emit (word, sum)

```

The MapReduce framework is a distributed implementation of a recursive algorithm. For example, we can sum up the elements of a list of integers as below.

```

1 def sum (xs: List[Int], z: Int = 0) : Int = xs match
2   case Nil => z
3   case y::ys => sum (ys, z + y)
4
5 val xs = List(11, 21, 31)
6 sum (xs)

```

The algorithm takes a list of (remaining) elements and a partial result, and aggregates the partial result with the result of processing a single element, until no more elements are left to process. It computes the sum in a *forward fashion*, passing the aggregated result to the next step. In a *backwards fashion* as below, the aggregation is postponed until all elements are processed.

```

1 def sum (xs: List[Int], z: Int = 0) : Int = xs match
2   case Nil => z
3   case y::ys => y + sum (ys, z)
4
5 val xs = List(11, 21, 31)
6 sum (xs)

```

Both have in common that they are using an accumulator. Scala has builtin fold operations that allow us to traverse collections while accumulating results. Operation `foldLeft` performs accumulation in a forward fashion, `foldRight` in backwards fashion. `foldLeft` is *tail recursive*, which means that the base case is the first element, the recursive call is on the tail, and the accumulator is applied to the head and the accumulated result. `foldRight` is *recursive into an argument*, which means that the base case is the last element, the recursive call is on the tail, and the accumulator is applied to the head and the result of the recursion. Folds are a universal concept that can be used to compute many different functions on lists, such as summing up elements, appending a list to another list, flattening, or reversing.

OPTION TYPES

Option types are a principled approach to missing data. `Option[T]` resembles `List[T]` with a length of at most 1. Java, for a long time, emphasized programming with exceptions in order to report anything that deviates from a successful result.

```

1 // the Java way
2 def getDirs1 (dirName : String) : List[java.io.File] =
3   val dir = new java.io.File (dirName)
4   val xs = dir.listFiles
5   if xs == null then throw new java.io.FileNotFoundException
6   xs.nn.toList.map (_.nn).filter (_.isDirectory)

```

Exceptions, however, are best used to report exceptional circumstances (such as a broken network connection); missing data is quite an expected result. Programming with optionals allows us to document missing data and allows our users to appropriately react to it (as opposed to the Java habit of just passing on exceptions).

```

1 def getDirs2 (dirName : String) : Option[ List[ java.io.File ] ] =
2   val dir = new java.io.File (dirName)
3   val xs = dir.listFiles
4   if xs == null then return None
5   Some(xs.nn.toList.map ( _ .nn ). filter ( _ .isDirectory ))

```

With optionals, clients no longer suffer from the convoluted way of providing a result from multiple execution traces. An option is a type that may have some result or nothing. Scala knows several option types:

- None is the empty option
- Nil is the empty list
- null is a reference to nothing

Unit is not an option type, it only has a single value (always has nothing). Even though Scala has null, we often pretend it does not exist. More recent languages, such as Swift and Kotlin, identify None and null; these languages distinguish nullable and non-nullable types. Java also includes an optional, but its intended use is narrowed to library methods whose return types needed a clear way of communicating the absence of a result and null is overwhelmingly likely to cause errors.