Week 11: Lisp

We now discuss the principles of another functional language: Lisp.

Lisp was the first functional language to be implemented; it dates back to 1959/60.

The name is an acronym for List processor.

At the time, Lisp was designed to deal with data structures for symbolic computation, such as lists or trees. (The languages that existed at the time only dealt with arrays).

Lisp has been used to implement many large applications. Examples:

- Macsyma, the first computer algebra program
- AutoCAD, a popular CAD program

Lisp Variants

In its 40+ years of existence, Lisp has evolved, and many dialects exist today.

The three most widespread ones are:

- Common Lisp (commercial, big),
- Scheme (academic, clean),
- Elisp (emacs extension language)

We treat here only a subset of Scheme.

The subset we present is purely functional – there are no variables or assignments. (Real Lisp has these and several other things as well).

Peculiarities of Lisp

Compared to Scala, there are four main points where Lisp differs and which make it interesting to study:

- No syntax, programs are simply nested sequences of words.
- No (static) type system.
- Only one form of compound data: the cons cell (from which lists are built).
- Programs are lists as well, so they can be constructed, transformed and evaluated by other program.

As an example of a Lisp program, consider *faculty*:

```
((define (faculty n)
    (if (= n 0)
        1
        (* n (faculty (- n 1)))))
 (faculty 10))
```

Observations:

- A compound Lisp expression is a sequence of sub-expressions in parentheses. Such an expression is also called a *combination*.
- Every subexpression is either a single word (an *atom*) or a compound expression.
- Subexpressions are separated by space characters.
- The first subexpression in a compound expression denotes an *operator* (usually a function).
- The other expressions denote the *operands*.

Special Forms

Some combinations look like function applications but aren't. In particular:

(*define name expr*)	defines *name* as an alias of the result of *expr* in the program that follows.
(*lambda* (*params*) *expr*)	the unnamed function that takes parameters *params* and returns *expr*, analogous to Scala's (*params* ⇒ *expr*).
(***if*** *cond expr1 expr2*)	returns result of *expr1* if *cond* evaluates to true, or the result of *expr2* otherwise.

Combinations like these are called *special forms*.

Function Applications

A combination $(op\ od_1\ ...\ od_n\ )$ which is not a special form is treated as a function application.

It is evaluated by applying the result of evaluating the operator op to the results of evaluating the operands od_1, ..., od_n.

This is essentially all. It is hard to imagine a useful programming language with fewer rules.

In fact, Lisp is so simple because it was meant as an intermediate language for computers, not humans.

It was originally planned to add a fancier Algol-like syntax on top of it, which would be translated to the intermediate form by a preprocessor.

However, humans got used to Lisp syntax rather quickly (at least some of them), and started to appreciate its advantages.

So the fancier syntax was never introduced.

Lisp Data

Data in Lisp are numbers, strings, symbols, and lists.

- Numbers are either floating point or integers. Integers have arbitrary size (no overflow!).
- Strings are as in Java.
- Symbols are simple sequences of characters without enclosing apostrophes. Examples:

  $x \quad head \quad + \quad \textbf{null}? \quad is{-}empty? \quad set!$

  Symbols are evaluated by looking up the value of a definition of the symbol in some environment.

There is no separate type of booleans. Instead, false is represented as the number *0*, whereas any value different from *0* is interpreted as true.

Lists in Lisp

Lists are written like combinations, e.g.

(1 2 3)
(1.0 "hello" (1 2 3))

Note that lists are heterogeneous; they can have elements of different types.

Note also that we may not evaluate a list like that, since its first element is not a function.

To prevent lists from being evaluated, we use the special form *quote*.

(quote (1 2 3))

The argument of *quote* is returned as the result without being evaluated itself. *quote* can be abbreviated with the special character '.

'(1 2 3)

Lists Internals

As in Scala, Lisp's list notation is just syntactic sugar.

Internally, lists are formed from:

- The empty list, written *nil*.
- Pairs of head x and tail y, written $(cons\ x\ y)$.

The list (or combination) $(x_1\ \ldots\ x_n)$ is represented as

$$cons\,(x_1,\ \ldots,\ cons\,(x_n,\ nil)\ \ldots\ )$$

Lists are accessed using the three operations

- $(\textbf{\textit{null?}}\ x)$ returns true if the list is empty.
- $(car\ x)$ returns the head of the list x.
- $(cdr\ x)$ returns the tail of the list x.

The names *car* and *cdr* point back to the IBM 704, the first computer on which Lisp was implemented.

On this computer, *CAR* meant "contents of address register" and *CDR* meant "contents of decrement register".

These two registers stored the two halves of a *cons* cell in the IBM 704 implementation.

In comparison to Scala,

cons	$\sim$	::
nil	$\sim$	*Nil*
null?	$\sim$	*isEmpty*
car	$\sim$	*head*
cdr	$\sim$	*tail*

Lists and Functions

Because Lisp has no static type system, we can represent lists using just functions, and a single symbol *none*. Here's how:

$$
\begin{aligned}
\mathit{nil} &= (\mathit{lambda}\ (k)\ (k\ \text{'}\mathit{none}\ \text{'}\mathit{none})) \\
(\mathit{cons}\ x\ y) &= (\mathit{lambda}\ (k)\ (k\ x\ y)) \\
(\mathit{car}\ l) &= (l\ (\mathit{lambda}\ (x\ y)\ x)) \\
(\mathit{cdr}\ l) &= (l\ (\mathit{lambda}\ (x\ y)\ y)) \\
(\mathbf{null?}\ l) &= (l\ (\mathit{lambda}\ (x\ y)\ (=\ x\ \text{'}\mathit{none})))
\end{aligned}
$$

The idea is that a cons-cell can be represented as a function which takes another function k as parameter.

The k function is intended to decompose the list.

The *cons* function simply applies k to its arguments.

Then, *car* and *cdr* simply apply the *cons* function to the appropriate decomposition functions.

This construction shows that in principle all data can be constructed from pure functions.

But in practice, a *cons* cell is represented by a pair of pointers.

An example

Here is the definition and usage of *map* in Lisp:

```
((define (map f xs)
   (if (null? xs)
       '()
       (cons (f (car xs)) (map f (cdr xs)))))
 (map (lambda (x) (* x x)) '(1 2 3)))
```

What is the result of evaluating this expression?

Interpeters

Lisp is so simple that it is an ideal vehicle for studying program evaluation rules.

In particular, it is fairly easy to write an interpreter for Lisp.

This is what we are going to do now.

More precisely, we will implement Scheme−−, a small but complete subset of Scheme, using Scala as the implementation language.

Defining new languages is something we did often in this course. Examples were:

- A language for arithmetic expressions,
- A language for digital circuits,
- A language for constraints.

All these languages were implemented as libraries of Scala functions.

Two things are new now:

- We will implement a full-blown programming language, which itself can compute any algorithm.
- The language will have an external syntax which looks like Lisp, not Scala.

  That syntax will be mapped to an internal data structure in Scala by a parser program.

The implementation will proceed in three stages.

1. Define internal representation of Scheme−− programs.
2. Define a translator from text strings to internal representations.
3. Define an interpreter program that evaluates internal representations.

Internal Representation of Scheme−−

Our internal representations of Scheme−− closely follow the data structures used in Scheme.

There is an abstract class *Data* which represents Scheme data.

```
abstract class Data;
```

There are subclasses to represent numbers, strings, symbols, and lists:

```
case class NUM(x: Int)                    extends Data;
case class STR(x: String)                 extends Data;
case class SYM(name: String)              extends Data;
case class NIL                            extends Data;
case class CONS(car: Data, cdr: Data)     extends Data;
```

Lists are built from *NIL* and *CONS*.

We also provide a set of overloaded *list* functions, that construct lists of different lengths:

```
def list ( ): Data = NIL;
def list (x1: Data): Data = CONS (x1, NIL);
def list (x1: Data, x2: Data): Data = CONS (x1, CONS (x2, NIL) );
...
```

and so on up to *list* (*x1*, ..., *x10*).

Example: The *faculty* function is written in Scheme−− as:

```
(define faculty (lambda (n)
   if (= n 0)
        1
        (* n (faculty (- n 1))))))
```

Its internal representation is constructed by:

```
list(SYM("define"), SYM("faculty"), list(SYM("lambda"), list(SYM("n"))),
   list(SYM("if"),
      list(SYM("="), SYM("n"), NUM(0)),
      NUM(1),
      list(SYM("*"),
         SYM("n"),
         list(SYM("faculty"), list(SYM("-"), SYM("n"), NUM(1))))))))
```

From Strings to Internal Representations

We now develop a parser program that maps strings into internal representation.

The mapping proceeds in two stages.

- From a string to a sequence of words (or: tokens)
- From a sequence of words to a *Data* tree.

Here, a token is one of the following.

- A left or right parenthesis "(", ")".
- A sequence of characters not containing whitespace or parentheses.
- Tokens which are not parentheses need to be separated by whitespace, i.e. blanks, new lines, or tabulator characters.

Tokenizing

We will represent the sequence of words in a Lisp program as an iterator.

The iterator is defined by the following class:

```
class LispTokenizer(s: String) extends Iterator[String] with {
    private var i = 0;
    private def isDelimiter(ch: Char) = ch ≤ ' ' || ch == '(' || ch == ')';
    def hasNext: Boolean = {
        while (i < s.length() && s.charAt(i) ≤ ' ') { i = i + 1 }
        i < s.length()
    }
    ...
```

Notes:

- The iterator keeps a private variable i, denoting the index of the next character to be read.

- The *hasNext* function skips any whitespace in front of the next token (if one exists). It uses the *charAt* method in Java's *String* class to access characters in the string.

```
...
def next : String =
    if (hasNext) {
        val start = i;
        var ch = s.charAt(i); i = i + 1;
        if (ch == '(') "("
        else if (ch == ')') ")"
        else {
            while (i < s.length() && !isDelimiter(s.charAt(i))) { i = i + 1 }
            s.substring(start, i)
        }
    } else error("premature end of string")
}
```

- The *next* function returns the next token. It uses the *substring* method in class *String*.

Parsing and Tree Construction

Given the extremely simple syntax of Lisp, it is possible to parse it without any advanced parsing techniques.

Here is how it is done.

```
def string2lisp(s: String): Data = {
    val it = LispTokenizer(s);
    def parseExpr(token: String): Data = {
        if (token == "(") parseList
        else if (token == ")") error("unbalanced parentheses")
        else if (Character.isDigit(token.charAt(0)))
            NUM(Integer.parseInt(token))
        else SYM(token)
    }
```

```
    def parseList: Data = {
        val token = it.next;
        if (token == ")") NIL else CONS(parseExpr(token), parseList)
    }
    parseExpr(it.next)
}
```

Notes:

- The *string2list* function converts a string into a *Data* expression tree.
- It first defines a *LispTokenizer* iterator *it*.
- This iterator is used by the two functions *parseExpr* and *parseList*, which call each other recursively.
- *parseExpr* parses a single expression.
- *parseList* parses a list of expressions in parentheses.

Now, if we write in the Scala interpreter:

```
string2lisp("(lambda (x) (+ (* x x) 1))")
```

we get (without the indentation):

```
CONS( SYM(lambda),
CONS( CONS( SYM(x), NIL),
CONS( CONS( SYM(+),
          CONS( CONS( SYM(*),
                    CONS( SYM(x),
                    CONS( SYM(x),
                    NIL))),
          CONS( NUM(1),
          NIL))),
NIL)))
```

Exercise: Add a method *asString* to the *Data* data constructors which prints Lisp expressions in Lisp form. I.e.

$$string2lisp(\text{"}(lambda\ (x)\ (+\ (*\ x\ x)\ 1))\text{"}).asString$$

should yield

$$(lambda\ (x)\ (+\ (*\ x\ x)\ 1))$$

Special Forms

Our Lisp interpreter will only be able to implement a single expression.

That expression can have the following special forms:

1.

$$(\textbf{\textit{val}}\ x\ expr_1\ \ expr_2\ )$$

Evaluates $expr_1$, binds the results to x, and then evaluates $expr_2$. This is analogous to Scala's $\textbf{\textit{val}}\ x = expr_1;\ expr_2$.

2.

$$(\textbf{\textit{def}}\ x\ expr_1\ \ expr_2\ )$$

Binds x to $expr_1$, and then evaluates $expr_2$. $expr_1$ is evaluated each time x is used. This is analogous to Scala's $\textbf{\textit{def}}\ x = expr_1;\ expr_2$.

Real Lisps have instead a range of binders called *define* for the toplevel and ***let***, ***let****∗ and *letrec* inside combinations.

Lisp's *define* and *letrec* correspond roughly to ***def*** and Lisp's ***let*** corresponds roughtly to ***val***, but their syntax is more complicated.

3.

$$(\mathit{lambda}\ (p_1\ \ldots\ p_n\,)\ \mathit{expr}\,)$$

Defines an anonymous function with parameters p_1, ..., p_n and body *expr*.

4.

$$(\mathit{quote}\ \mathit{expr}\,)$$

returns *expr* without evaluating it.

5.

$$(\boldsymbol{if}\ \mathit{cond}\ \mathit{expr}_1\ \mathit{expr}_2\,)$$

The usual if-then-else construct.

Removing Syntactic Sugar

Other forms can be mapped to these by transforming the internal representation of Lisp data.

One can write a function *normalize* that eliminates other special forms in trees.

For instance, Lisp supports the special form

$(and\ x\ y)$
$(or\ x\ y)$

for short-circuited and and or. The *normalize* procedure would map these to if-then-else expressions as follows.

```
def normalize(expr: Data) = expr match {
    case CONS(SYM("and"), CONS(x, CONS(y, NIL))) ⇒
        list(SYM("if"), x, y, NUM(0))
    case CONS(SYM("or"), CONS(x, CONS(y, NIL))) ⇒
        list(SYM("if"), x, NUM(1), y)
    ... // further simplifications go here
    case NIL ⇒ NIL
    case CONS(x, y) ⇒ CONS(normalize x, normalize y)
}
```

Summary

Lisp is quite an unusual language.

- It does away with elaborate syntax and static type systems.
- It reduces all compound data to lists.
- It identifies programs and data.

These points have advantages as well as disadvantages.

- No syntax:
  + easy to learn, − hard to read.
- No static type system.
  + flexible, − easy to make mistakes.
- Only lists as compound data:
  + flexible, − poor data abstraction.
- Programs are data:
  + powerful, − hard to keep safe.