

Closure conversion

Michel Schinz – parts based on slides by Xavier Leroy
2007-04-27

Higher-order functions

Higher-order function

A **higher-order function (HOF)** is a function that either:

- takes another function as argument, or
- returns a function.

Many languages offer higher-order functions, but not all provide the same power...

HOFs in C

In C, it is possible to pass a function as an argument, and to return a function as a result.

However, C functions cannot be nested: they must all appear at the top level. This severely restricts their usefulness, but greatly simplifies their implementation – they can be represented as simple code pointers.

HOFs in functional languages

In functional languages – Scala, Scheme, OCaml, etc. – functions can be nested, and they can survive the scope that defined them.

This is very powerful as it permits the definition of functions that return “new” functions – e.g. functional composition.

However, as we will see, it also complicates the representation of functions, as simple code pointers are no longer sufficient.

HOF example

To illustrate the issues related to the representation of functions in a functional language, we will use the following Scheme example:

```
(define make-adder
  (lambda (x)
    (lambda (y) (+ x y))))

(define increment (make-adder 1))
(increment 41) ⇒ 42

(define decrement (make-adder -1))
(decrement 42) ⇒ 41
```

Representing adder functions

To represent the functions returned by `make-adder`, we basically have two choices:

1. Keep the code pointer representation for functions. However, that implies run-time code generation, as each function returned by `make-adder` is different!
2. Find another representation for functions, which does not depend on run-time code generation.

Closures

Closures

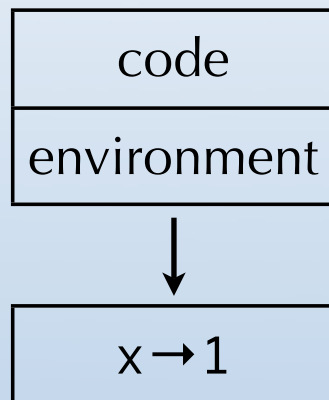

To adequately represent the functions returned by `make-adder`, their code pointer must be augmented with the value of `x`.

Such a combination of a **code pointer** and an **environment** giving the values of the free variable(s) – here `x` – is called a **closure**.



The name refers to the fact that the pair (code pointer, environment) is self-contained.

Closure

(make-adder 1)



(make-adder -1)



shared code

compiled code for:
(lambda (x)
 (lambda (y)
 (+ x y)))

The code of a closure must be evaluated in its environment, so that x is “known”.

Introducing closures

Using closures instead of function pointers to represent functions changes the way they are manipulated at run time:

- function abstraction builds and returns a closure instead of a simple code pointer,
- function application extracts the code pointer from the closure, and invokes it with the environment as an additional argument.

Representing closures

During function application, nothing is known about the closure being called – it can be any closure in the program.




The code pointer must therefore be at a known and constant location so that it can be extracted.

The *content* of the environment, however, is not used during application itself: it will only be accessed by the function body. This provides some liberty to represent it.

Flat closures

In flat (or one-block) closures, the environment is “inlined” into the closure itself, instead of being referred from it. The closure plays the role of the environment.

(make-adder 1)



Recursive closures

Recursive functions need access to their own closure. For example:

```
(define f  
  (lambda (l) ... (map f l) ...))
```

Several techniques can be used to give a closure access to itself:

- the closure – here `f` – can be treated as a free variable, and put in its own environment – leading to a cyclic closure,
- the closure can be rebuilt from scratch,
- with flat closures, the environment is the closure, and can be reused directly.

Mutually-recursive closures

Mutually-recursive functions all need access to the closures of all the functions in the definition.

For example, in the following program, `f` needs access to the closure of `g`, and the other way around:


```
(letrec ((f (lambda (l) ...(compose f g)...))
         (g (lambda (l) ...(compose g f)...)))
  ...)
```

Solutions:

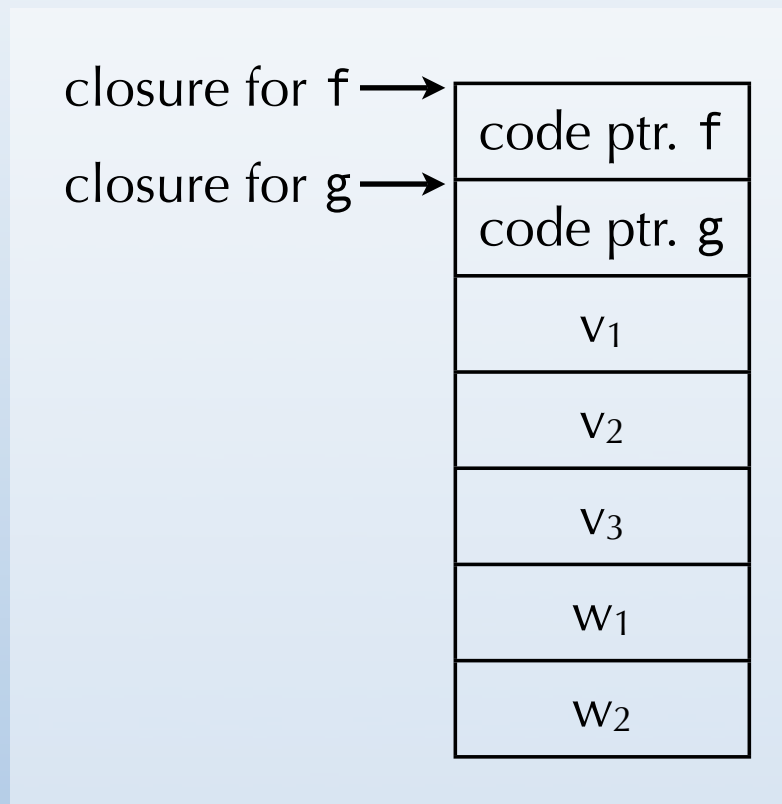
1. use cyclic closures, or
2. share a single closure with interior pointers.

Mutually-recursive closures

cyclic closures



shared closure



Compiling closures

Closure conversion

In a compiler, closures can be implemented by a simplification phase, called **closure conversion**.

Closure conversion transforms a program in which functions can be nested and have free variables into an equivalent one containing only top-level – and hence closed – functions.

The output of closure conversion is therefore a program in which functions can be represented as code pointers!

Closure conversion phases

Closure conversion can be split in two phases:

- the closing of functions, through the introduction of environments,
- the hoisting of nested, closed functions to the top level.

We will examine them later, but we first need to define the concept of free variable.

Free variables

The **free variables** of a function are the variables that are used but not defined in that function – *i.e.* they are defined in some enclosing scope.

Global variables are never considered free, since they are available everywhere.

Free variables example

Our adder example contains two functions, corresponding to the two occurrences of the `lambda` keyword:

```
(define make-adder  
  (lambda (x)  
    (lambda (y) (+ x y))))
```

The **outer one** does not have any free variable: it is a **closed function**, like all top-level functions. The **inner one** has a single free variable: `x`.

Closing functions


Functions are closed by adding a parameter representing the environment, and using it in the function's body to access free variables.

Function abstraction and application must of course be adapted accordingly:

- abstraction must create and initialise the closure and its environment,
- application must extract the environment and pass it as an additional parameter.

Closing example

```
(define make-adder  
  (lambda (x)  
    (lambda (y) (+ x y))))
```



```
(define make-adder  
  (vector (lambda (env1 x)  
            (vector (lambda (env2 y)  
                      (+ (vector-ref env2 1) y))  
                      x))))))
```

closure for
make-adder

closure
for anonymous
adder

Hoisting functions


Once they are closed, nested anonymous functions can easily be hoisted to the top level and given an arbitrary name.

The original occurrence of the nested function is simply replaced by that name.

After hoisting, all functions appearing in the program are at the top-level, and are of course closed. Therefore, they can be represented by simple code pointers, as in C.

Hoisting example

```
(define make-adder  
  (vector (lambda (env1 x)  
            (vector (lambda (env2 y)  
                      (+ (vector-ref env2 1) y))  
                    x))))))
```



```
(define lambda3  
  (lambda (env2 y)  
    (+ (vector-ref env2 1) y)))  
(define lambda4  
  (lambda (env1 x)  
    (vector lambda3 x)))  
(define make-adder  
  (vector lambda4))
```

Closure conversion for minischeme

Minischeme closure conversion

As we have seen, closure conversion can be performed by first closing functions, and then hoisting nested functions to the top level.

We will look in detail at the closing part for minischeme, which we will specify as a function C mapping potentially-open terms to closed ones.

For that, we first need to define a function F mapping a term to the set of its free variables.

Note: to simplify presentation, we assume in the following slides that all variables in a program have a unique name.

Minischeme free variables

$$F[(\text{define name value})] = \emptyset$$

$$F[(\text{lambda } (v_1 \dots) \text{body}_1 \dots)] = \\ (F[\text{body}_1] \cup F[\text{body}_2] \cup \dots) \setminus \{v_1, \dots\}$$

$$F[(\text{let } ((v_1 e_1) \dots) \text{body}_1 \dots)] = \\ (F[e_1] \cup \dots \cup F[\text{body}_1] \cup \dots) \setminus \{v_1, \dots\}$$

$$F[(\text{if } e_1 e_2 e_3)] = F[e_1] \cup F[e_2] \cup F[e_3]$$

$$F[(e_1 e_2 \dots)] = F[e_1] \cup F[e_2] \cup \dots$$

$$F[v] \text{ when } v \text{ is local} = \{v\}$$

$$F[v] \text{ when } v \text{ is global or a primitive} = \emptyset$$

Closing minischeme functions

Closing minischeme constructs that do not deal with functions or variables is trivial:

$$c[(\text{define name value})] = \\ (\text{define name } c[\text{value}])$$
$$c[(\text{let } ((v_1 e_1) \dots) \text{body}_1 \dots)] = \\ (\text{let } ((v_1 c[e_1]) \dots) c[\text{body}_1] \dots)$$
$$c[(\text{if } e_1 e_2 e_3)] = \\ (\text{if } c[e_1] c[e_2] c[e_3])$$
$$c[n] \text{ where } n \text{ is a number} = \\ n$$

Closing minischeme functions

Abstraction is closed by creating and returning the closure, represented as a vector:

$$c[(\text{lambda } (v_1 \dots) \text{body}_1 \dots)] = \\ (\text{vector } (\text{lambda } (\underline{\text{env}} v_1 \dots) \epsilon[C[\text{body}_1], F, \underline{\text{env}}] \dots) \\ F@1 F@2 \dots)$$

where

- $\epsilon[t, f, e]$ transforms t by replacing all occurrences of the variables of f by accesses to corresponding slots in the environment e .
- $F = F[(\text{lambda } (v_1 \dots) \text{body}_1 \dots)]$

Closing minischeme functions

Finally, application extracts the code pointer from the closure, and invokes it with the closure itself as the first argument, followed by the other arguments:

$C[(e_1 e_2 \dots)]$ when e_1 is not a primitive =
`(let ((closure C[e1]))
 ((vector-ref closure 0) closure C[e2] ...))`

$C[(e_1 e_2 \dots)]$ when e_1 is a primitive =
`(e1 C[e2] ...)`

Hoisting minischeme functions

Hoisting consists in lifting closed nested functions to the top-level, naming them in the process.

Like closing, hoisting can be specified as a function, say H , mapping a term that potentially contains nested functions, to a new version of that term without nested functions – plus a list of additional definitions.

The definition of that function is left as an exercise.

Closures and objects

Closures and objects

There is a strong similarity between closures and objects: closures can be seen as objects with a single method – containing the code of the closure – and a set of fields – the environment.

In Java, the ability to define nested classes can be used to simulate closures, but the syntax is too heavyweight to be used often.

In Scala, a special syntax exists for anonymous functions, which are translated to nested classes.

Closures in Scala

To see how closures are handled in Scala, we will look at how the compiler translates the Scala equivalent of the make-adder function:

```
def makeAdder(x: Int): Int=>Int =  
  { y: Int => x+y }  
val increment = makeAdder(1)  
increment(41)
```

Closures in Scala

In a first phase, the anonymous function is turned into an anonymous class of type `Function1` – the type of functions with one argument. This class is equipped with a single `apply` method containing the code of the anonymous function.

```
def makeAdder(x: Int): Function1[Int,Int]=  
  new Function1[Int,Int] {  
    def apply(y: Int): Int = x+y  
  }  
val increment = makeAdder(1)  
increment.apply(41)
```

Closures in Scala

In a second phase, the anonymous class is named.

```
def makeAdder(x: Int):Function1[Int,Int]={
  class Anon extends Object
    with Function1[Int,Int] {
    def apply(y: Int): Int = x+y
  }
  new Anon
}
val increment = makeAdder(1)
increment.apply(41)
```

Closures in Scala

In a third phase, the Anon class is closed and hoisted to the top level.

```
class Anon(x:Int) extends Object
    with Function1[Int,Int]{
  def apply(y: Int): Int = x+y
}
def makeAdder(x: Int):Function1[Int,Int]={
  new Anon(x)
}
val increment = makeAdder(1)
increment.apply(41)
```

Closures in Scala

Finally, the constructor of Anon is made explicit.

```
class Anon extends Object
    with Function1[Int,Int] {
  private var x: Int = _;
  def this(x0: Int) { this.x = x0 }
  def apply(y: Int): Int = x+y
}
def makeAdder(x: Int):Function1[Int,Int]={
  new Anon(x)
}
val increment = makeAdder(1)
increment.apply(41)
```

Summary

In C, all functions have to be at the top level, and can therefore be represented as code pointers.

Functional languages allow functions to be nested and to survive the scope that created them. They have to be represented by a closure, which pairs a code pointer with an environment giving the values of the code's free variables.

Closures can be implemented by a program transformation called closure conversion.