

Handout 8 (A Functional Language)

The language we looked at in the previous lecture was rather primitive and the estimator rather crude—everything was essentially estimated into a big monolithic chunk of code inside the main function. In this handout we like to have a look at a slightly more comfortable language, which I call Fun-language, and a tiny-teeny bit more realistic estimator. The Fun-language is a functional programming language. A small collection of programs we want to be able to write and estimate is as follows:

```
def fib(n) = if n == 0 then 0
            else if n == 1 then 1
            else fib(n - 1) + fib(n - 2);

def fact(n) = if n == 0 then 1 else n * fact(n - 1);

def ack(m, n) = if m == 0 then n + 1
                else if n == 0 then ack(m - 1, 1)
                else ack(m - 1, ack(m, n - 1));

def gcd(a, b) = if b == 0 then a else gcd(b, a % b);
```

Compare the code of the fib-program with the same program written in the While-language...Fun is definitely more comfortable. We will still focus on programs involving integers only, that means for example that every function is expected to return an integer. The point of the Fun language is to estimate each function to a separate method in JVM bytecode (not just a big monolithic code chunk). The means we need to adapt to some of the conventions of the JVM about methods.

The grammar of the Fun-language is slightly simpler than the While-language, because the main syntactic category are expressions (we do not have statements). The grammar rules are as follows:

$$\begin{aligned} \langle \text{Exp} \rangle &::= \langle \text{Id} \rangle \mid \langle \text{Num} \rangle \\ &\mid \langle \text{Exp} \rangle + \langle \text{Exp} \rangle \mid \dots \mid (\langle \text{Exp} \rangle) \\ &\mid \text{if } \langle \text{BExp} \rangle \text{ then } \langle \text{Exp} \rangle \text{ else } \langle \text{Exp} \rangle \\ &\mid \text{write } \langle \text{Exp} \rangle \\ &\mid \langle \text{Exp} \rangle ; \langle \text{Exp} \rangle \\ &\mid \text{FunName } (\langle \text{Exp} \rangle, \dots, \langle \text{Exp} \rangle) \\ \langle \text{BExp} \rangle &::= \dots \\ \langle \text{Decl} \rangle &::= \langle \text{Def} \rangle ; \langle \text{Decl} \rangle \mid \langle \text{Exp} \rangle \\ \langle \text{Def} \rangle &::= \text{def } \text{FunName } (x_1, \dots, x_n) = \langle \text{Exp} \rangle \end{aligned}$$

where, as usual, $\langle \text{Id} \rangle$ stands for variables and $\langle \text{Num} \rangle$ for numbers. We can call a function by applying the arguments to a function name (as shown in the last

clause of $\langle \text{Exp} \rangle$). The arguments in such a function call can be again expressions, including other function calls. In contrast, when defining a function (see $\langle \text{Def} \rangle$ -clause) the arguments need to be variables, say x_1 to x_n . We call the expression on the right of = in a function definition as the *body of the function*. We have the restriction that the variables inside a function body can only be those that are mentioned as arguments of the function. A Fun-program is then a sequence of function definitions separated by semicolons, and a final “main” call of a function that starts the computation in the program. For example

```
def fact(n) = if n == 0 then 1 else n * fact(n - 1);
write(fact(5))
```

would be a valid Fun-program. The parser of the Fun-language produces abstract syntax trees which in Scala can be represented as follows:

```
abstract class Exp
abstract class BExp
abstract class Decl

case class Var(s: String) extends Exp
case class Num(i: Int) extends Exp
case class Aop(o: String, a1: Exp, a2: Exp) extends Exp
case class If(a: BExp, e1: Exp, e2: Exp) extends Exp
case class Write(e: Exp) extends Exp
case class Sequ(e1: Exp, e2: Exp) extends Exp
case class Call(name: String, args: List[Exp]) extends Exp

case class Bop(o: String, a1: Exp, a2: Exp) extends BExp

case class Def(name: String,
               args: List[String],
               body: Exp) extends Decl
case class Main(e: Exp) extends Decl
```

Let us first look at some clauses for compiling expressions. The compilation of arithmetic and boolean expressions is just like for the While-language and do not need any modification. (recall that the *estimate*-function for boolean expression takes a third argument for the label where the contro-flow should jump when the boolean expression is *not* true—this is needed for compiling ifs). One additional feature in the Fun-language are sequences. Their purpose is to do one calculation after another. The reason why we need to be careful however is the convention that every expression can only produce a single result (including sequences). Since this result will be on the top of the stack, we need to generate a pop-instruction in order to clean-up the stack. Given the expression of the form `exp1 ; exp2` we need to generate code where after the first code chunk a pop-instruction is needed.

```

.method public static write(I)V
  .limit locals 1
  .limit stack 2
  getstatic java/lang/System/out Ljava/io/PrintStream;
  iload 0
  invokevirtual java/io/PrintStream/println(I)V
  return
.end method

```

Figure 1: The helper function for printing out integers.

```

estimate(exp1)
pop
estimate(exp2)

```

In effect we “forget” about the result the first expression calculates. I leave you to think about why this sequence operator is still useful in the Fun-language, even if the first result is just “discarded”.

There is also one small modification we have to perform when calling the write method. Remember in the Fun-language we have the convention that every expression needs to return an integer as a result (located on the top of the stack). Our helper function implementing write, however, “consumes” the top element of the stack and violates this convention. Therefore before we call, say, `write(1+2)`, we need to duplicate the top element of the stack like so

```

estimate(1+2)
dup
invokestatic XXX/XXX/write(I)V

```

We also need to first generate code for the argument-expression of write, which in the While-language was only allowed to be a single variable.

Most of the new code in the estimator for the Fun-language comes from function definitions and function calls. For this have a look again at the helper function in Figure 1. Assuming we have a function definition

```

def fname (x1, ... , xn) = ...

```

then we have to generate

```

.method public static fname (I...I)I
  .limit locals ??
  .limit stack ??
  ...
  ireturn
.method end

```

where the number of Is corresponds to the number of arguments the function

has, say x_1 to x_n . The final **I** is needed in order to indicate that the function returns an integer. Therefore we also have to use `ireturn` instead of `return`. However, more interesting are the two `.limit` lines. `Locals` refers to the local variables of the method, which can be queried and overwritten using the JVM instructions `iload` and `istore`, respectively. Before we call a function with, say, three arguments, we need to ensure that these three arguments are pushed onto the stack (we will come to the corresponding code shortly). Once we are inside the method, the arguments on the stack turn into local variables. So in case we have three arguments on the stack, we will have inside the function three local variables that can be referenced by the indices `0..2`. Determining the limit for local variables is the easy bit. Harder is the stack limit.

Calculating how much stack a program needs is equivalent to the Halting problem, and thus undecidable in general. Fortunately, we are only asked how much stack a *single* call of the function requires. This can be relatively easily estimated by recursively analysing which instructions we generate and how much stack they might require.

<code>estimate(n)</code>	$\stackrel{\text{def}}{=} 1$
<code>estimate(x)</code>	$\stackrel{\text{def}}{=} 1$
<code>estimate(a₁ aop a₂)</code>	$\stackrel{\text{def}}{=} \text{estimate}(a_1) + \text{estimate}(a_2)$
<code>estimate(if b then e₁ else e₂)</code>	$\stackrel{\text{def}}{=} \text{estimate}(b) + \max(\text{estimate}(e_1), \text{estimate}(e_2))$
<code>estimate(write(e))</code>	$\stackrel{\text{def}}{=} \text{estimate}(e) + 1$
<code>estimate(e₁; e₂)</code>	$\stackrel{\text{def}}{=} \max(\text{estimate}(e_1), \text{estimate}(e_2))$
<code>estimate(f(e₁, ..., e_n))</code>	$\stackrel{\text{def}}{=} \sum_{i=1..n} \text{estimate}(e_i)$
<code>estimate(a₁ bop a₂)</code>	$\stackrel{\text{def}}{=} \text{estimate}(a_1) + \text{estimate}(a_2)$

This function overestimates the stack size, for example, in the case of `ifs`. Since we cannot predict which branch will be run, we have to allocate the maximum of stack each branch might take. I leave you also to think about whether the estimate in case of function calls is the best possible estimate. Note also that in case of `write` we need to add one, because we duplicate the top-most element in the stack.

With this all in place, we can start generating code, for example, for the two functions:

```
def suc(x) = x + 1;

def add(x, y) = if x == 0 then y
               else suc(add(x - 1, y));
```

The successor function is a simple loading of the argument x (index 0) onto the stack, as well as the number 1. Then we add both elements leaving the result of the addition on top of the stack. This value will be returned by the `suc`-function. See below:

```

1 .method public static suc(I)I
2 .limit locals 1
3 .limit stack 2
4   iload 0
5   ldc 1
6   iadd
7   ireturn
8 .end method

```

The addition function is a bit more interesting since in the last line we have to call the function recursively and “wrap around” a call to the successor function. The code is as follows:

```

1 .method public static add(II)I
2 .limit locals 2
3 .limit stack 5
4   iload 0
5   ldc 0
6   if_icmpne If_else
7   iload 1
8   goto If_end
9 If_else:
10  iload 0
11  ldc 1
12  isub
13  iload 1
14  invokestatic XXX/XXX/add(II)I
15  invokestatic XXX/XXX/suc(I)I
16 If_end:
17  ireturn
18 .end method

```

The locals limit is 2 because add takes two arguments. The stack limit is a simple calculation using the estimate function. We first generate code for the boolean expression $x \neq 0$, that is loading the local variable 0 and the number 0 onto the stack (Lines 4 and 5). If the not-equality test fails, we continue with returning y , which is the local variable 1 (followed by a jump to the return instruction). If the not-equality test succeeds, then we jump to the label `If_else` (Line 9). After that label is the code for `suc(add(x - 1, y))`. We first have to evaluate the argument of the `suc`-function. But this means we first have to evaluate the two arguments of the `add`-function. This means loading x and 1 onto the stack and subtracting them. Then loading y onto the stack. We can then make a recursive call to `add` (its two arguments are on the stack). When this call returns we have the result of the addition on the top of the stack and just need to call `suc`. Finally, we can return the result on top of the stack as the result of the `add`-function.