

Handout 9 (LLVM, SSA and CPS)

Reflecting on our tiny compiler targetting the JVM, the code generation part was actually not so hard, no? Pretty much just some post-traversal of the abstract syntax tree, yes? One of the main reason for this ease is that the JVM is a stack-based virtual machine and it is therefore not hard to translate arithmetic expressions into a sequence of instructions manipulating the stack. The problem is that “real” CPUs, although supporting stack operations, are not really designed to be *stack machines*. The design of CPUs is more like, here is a chunk of memory—compiler, or better compiler writers, do something with it. Consequently, modern compilers need to go the extra mile in order to generate code that is much easier and faster to process by CPUs.

Another reason why it makes sense to go the extra mile is that stack instructions are very difficult to optimise—you cannot just re-arrange instructions without messing about with what is calculated on the stack. Also it is hard to find out if all the calculations on the stack are actually necessary and not by chance dead code. The JVM has for all this sophisticated machinery to make such “high-level” code still run fast, but let’s say that for the sake of argument we do not want to rely on it. We want to generate fast code ourselves. This means we have to work around the intricacies of what instructions CPUs can actually process. To make this all tractable for this module, we target the LLVM Intermediate Language. In this way we can take advantage of the tools coming with LLVM. For example we do not have to worry about things like register allocations.

LLVM¹ is a beautiful example that projects from Academia can make a difference in the world. LLVM started in 2000 as a project by two researchers at the University of Illinois at Urbana-Champaign. At the time the behemoth of compilers was gcc with its myriad of front-ends for other languages (e.g. Fortran, Ada, Go, Objective-C, Pascal etc). The problem was that gcc morphed over time into a monolithic gigantic piece of m...ehm software, which you could not mess about in an afternoon. In contrast, LLVM is designed to be a modular suite of tools with which you could play around easily and try out something new. LLVM became a big player once Apple hired one of the original developers (I cannot remember the reason why Apple did not want to use gcc, but maybe they were also just disgusted by its big monolithic codebase). Anyway, LLVM is now the big player and gcc is more or less legacy. This does not mean that programming languages like C and C++ are dying out any time soon—they are nicely supported by LLVM.

Targetting the LLVM Intermediate Language, or Intermediate Representation (short LLVM-IR), also means we can profit from the very modular structure of the LLVM compiler and let for example the compiler generate code for X86, or ARM etc. That means we can be agnostic about where our code actually runs.

© Christian Urban, King’s College London, 2019

¹<http://llvm.org>

However, what we have to do is to generate code in *Static Single-Assignment* format (short SSA), because that is what the LLVM-IR expects from us. LLVM-IR is the intermediate format that LLVM uses for doing cool things, like targetting strange architectures, optimising code and allocating memory efficiently.

The idea behind the SSA format is to use very simple variable assignments where every variable is assigned only once. The assignments also need to be primitive in the sense that they can be just simple operations like addition, multiplication, jumps, comparisons and so on. An idealised snippet of a program in SSA is

```
x := 1
y := 2
z := x + y
```

where every variable is used only once (we could not write `x := x + y` in the last line for example). There are sophisticated algorithms for imperative languages, like C, that efficiently transform a high-level program into SSA format. But we can ignore them here. We want to compile a functional language and there things get much more interesting than just sophisticated. We will need to have a look at CPS translations, where the CPS stands for Continuation-Passing-Style—basically black programming art or abracadabra programming. So sit tight.

LLVM-IR

Before we start, lets first have a look at the *LLVM Intermediate Representation*. What is good about our simple Fun language is that it basically only contains expressions (be they arithmetic expressions or boolean expressions). The exception is function definitions. Luckily, for them we can use the mechanism of defining functions in LLVM-IR. For example the simple Fun program

```
def sqr(x) = x * x
```

can be compiled into the following LLVM-IR function:

```
define i32 @sqr(i32 %x) {
    %tmp = mul i32 %x, %x
    ret i32 %tmp
}
```

First to notice is that all variable names in the LLVM-IR are prefixed by %; function names need to be prefixed with @. Also, the LLVM-IR is a fully typed language. The `i32` type stands for a 32-bit integer. There are also types for 64-bit integers, chars (`i8`), floats, arrays and even pointer types. In teh code above, `sqr` takes an argument of type `i32` and produces a result of type `i32`. Each arithmetic operation, like addition or multiplication, are also prefixed with the type they operate on. Obviously these types need to match up... but since we have in our programs only integers, `i32` everywhere will do.

```

@.str = private constant [4 x i8] c"%d\0A\00"

declare i32 @printf(i8*, ...)

; print out an integer
define i32 @printInt(i32 %x) {
    %t0 = getelementptr [4 x i8], [4 x i8]* @.str, i32 0, i32 0
    call i32 @printf(i8* %t0, i32 %x)
    ret i32 %x
}

; square function
define i32 @sqr(i32 %x) {
    %tmp = mul i32 %x, %x
    ret i32 %tmp
}

; main
define i32 @main() {
    %1 = call i32 @sqr(i32 5)
    %2 = call i32 @printInt (i32 %1)
    ret i32 %1
}

```

Figure 1:

Conveniently, you can use the program `lli`, which comes with LLVM, to interpret programs written in the LLVM-IR. So you can easily check whether the code you produced actually works. To get a running program that does something interesting you need to add some boilerplate about printing out numbers and a main-function that is the entrypoint for the program (see Figure 1).

CPS-Translations

```
abstract class Exp extends Serializable
abstract class BExp extends Serializable
abstract class Decl extends Serializable

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

case class Call(name: String, args: List[Exp]) extends Exp
case class If(a: BExp, e1: Exp, e2: Exp) extends Exp
case class Write(e: Exp) extends Exp
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 Sequence(e1: Exp, e2: Exp) extends Exp
case class Bop(o: String, a1: Exp, a2: Exp) extends BExp
```

Figure 2: Abstract syntax trees for the Fun language.