

Machine Problem 2

CS 426 — Compiler Construction
Fall Semester 2022

Handed out: September 14, 2022. Due: October 5, 2022, 5 pm

In this assignment and the next (MP 2 and MP 3), you will implement the intermediate code generation phase of your compiler. In MP 2, you'll implement only the language features of COOL shown in the figure below.

1 COOL Language Subset for MP 2

```

program ::= class;
class ::= class Main { feature; }
feature ::= main() : Int { expr }
expr ::= ID <- expr
        | if expr then expr else expr fi
        | while expr loop expr pool
        | { expr;+ }
        | let [[ID : TYPE [ <- expr ]],+ in expr
        | expr + expr
        | expr - expr
        | expr * expr
        | expr / expr
        | ~ expr
        | expr < expr
        | expr <= expr
        | expr = expr
        | not expr
        | (expr)
        | ID
        | integer
        | true
        | false

```

Figure 1: Syntax for the subset of COOL to be implemented in MP 2.

Essentially, you are leaving out language features that require implementing classes, including methods and method calls, attributes, inheritance, constructors, **new** and **case**. The class **Main** is the only class, and it is assumed to have a single method **Main.main() : Int**. This method becomes a simple, global

LLVM function. (We have given you special-purpose code to translate class `Main` and method `main()`; you only need to translate its body.)

Note that the only types you must support are `Int` and `Bool`; in particular, `String` is not included because it requires objects, and `SELF_TYPE` is not needed in the absence of objects. To eliminate objects from the language, we need to make two small changes to the Cool typing rules:

1. You can assume that both branches of a conditional expression have the same type (both `Int` or both `Bool`). Therefore, the type of the whole expression will be either `Int` or `Bool`, and it will be the same as the type of the branches. In MP 3, you will have to handle the case of *different* types in the branches, that are merged into the “join” of the two types (see Section 7.5 of the Cool manual).
2. A loop expression has type `Int` and evaluates to the value 0. In MP 3, you must implement the rule that a loop expression evaluates to a `void` value of type `Object` (Cool manual, Section 7.6). However, for MP 2, `Object` is not supported.

The Cool runtime library is not used for MP 2. Also, you cannot use class `IO` to perform input-output. Instead, you can return a result from the function `main`.

The code generator makes use of the AST constructed in MP1 and static analysis performed by the program *semant*. Your code generator should produce LLVM assembly code that faithfully implements *any* correct Cool program matching the language subset described above. In particular, there is no static error checking in code generation—all erroneous Cool programs have been detected by the front-end phases of the compiler (except for errors that must be detected at run-time, such as divide-by-zero in MP 2).

This assignment gives you some flexibility in how exactly you generate LLVM code for individual Cool constructs. We will specify certain design choices in order to simplify the project (e.g., how to organize classes in MP 3, and how to return values to the environment in MP 2). You are responsible for other key design choices, including how to implement the basic built-in classes (`Int` and `Bool` in MP 2). Nevertheless, note that there are many key design goals to meet, and there are standard design approaches compilers use to meet these goals. We will discuss these approaches in class or in this handout.

2 LLVM Environment and Usage

See the handout, “**Introductory LLVM Exercise**” (under Resources on the webpage) for how to set up your environment to run the LLVM tools, and for a simple exercise to get familiar with the LLVM IR.

3 MP 2 Distribution

The file *mp2.tar.gz* is available on the course web site, under the Project page. The *mp2* directory contains a README file with documentation on the layout of the skeleton code and the overall structure of the code generator. Read this before you begin writing code. Also familiarize yourself with the *Tour of the COOL Support Code* document on the course web site, so that you can refer back to it as needed.

The subdirectory *mp2/src* contains the skeleton files for the code generation phase, which you will need to modify. You should not need to change any files in any other directories. The *mp2/src* directory contains:

- *cgen.cc*:
This file will contain your code generator. We have provided an implementation of some aspects of code generation; studying this code will help you write the rest of the code generator. It includes a call to code that will build an inheritance graph from the provided AST.
- *cgen.h*:
This file is the header for the code generator. You may add anything you like to this file. It also provides classes for implementing the inheritance graph.
- *cool-tree.handcode.h*, *stringtab.handcode.h* :
The former modifies the declarations of classes for the AST nodes. The latter modifies the declarations of classes *StrTable* and *IntTable*. You can add field declarations to these classes by editing these two files. The macros defined in these files are included into *cool-tree.h* and *stringtab.h* respectively when building your compiler.

So if you want to add a function *void addMe(int)* as pure virtual into for example the *Feature_class* and its implementation into all subclasses, you have to put a definition *virtual void addMe(int)=0* into *Feature_EXTRAS* and *void addMe(int);* into *Feature_SHARED_EXTRAS*. Take a look at *cool-tree.handcode.h* and *cool-tree.h* to fully understand this.

The definitions of the methods should be added to *cgen.cc*.
- *Makefile*:
The Makefile for the *src/* directory. It is similar to the Makefile for MP1 but has additional targets. You can modify this Makefile (for example, to change the *debug* flag) but this file will not be turned in, so your compiler must be able to build using the provided version.
- *value_printer.{cc,h}*, *operand.{cc,h}*: These files contain a small library for printing out LLVM bytecode assembly files, which you may use for the assignment. The provided *cgen.h* includes *value_printer.h*. You may also print your assembly directly to the output stream, if you prefer (although you'll have to make sure to get the syntax right). Using the internal API to construct the API as in-memory data structures and pretty-print the IR would require substantially more work, so we have chosen the simpler route of printing out LLVM bytecode assembly as text files.
- *coolrt.{c,h}*: These files provide a partial implementation of the COOL runtime library, for you to complete in MP 3. You will not be using them for MP 2.

To compile your code generator for MP 2 type *make cgen-1*. Like the Makefile for MP1, the Makefile will compile files from the *cool-support* directory and link the object files into your program.

Note that *cgen.h* and *cgen.cc* use conditional compilation directives (*ifdef* and *ifndef*) to build two different programs, depending on whether the symbol **MP3** is defined. We have set up the Makefile so that when you build *cgen-1*, **MP3** will not be defined, and these regions will not be compiled, nor will they appear in your binary program *cgen-1*. You should not add any code to these sections for MP 2 (you will need to, later, for MP 3).

The directory *reference-binaries* includes our binary tools for the previous phases of the COOL compiler (lexer, parser, semantic analyzer). It also contains our code generator for MP 2 called *cgen-1*, which you can use as a reference; it shows you what code to generate for test cases.

Read the document, *A Tour of the COOL Support Code*, available at the **Resources** link on the web site. You should also take a short look at the other files in the *cool-support/src*, and *cool-support/include*

directories. Especially *cgen_supp.cc* contains general support code for the code generator. You will find a number of handy functions here.

4 Testing the Code Generator

The directory *mp2/test-1/* provides a place for you to test your code generator for MP 2. *You should write your own test cases to test your compiler.* Use separate simple tests initially, e.g., a single constant and simple arithmetic with two constants, and then work your way up to more complex expressions. A couple of days before the due date, we'll make the handin script available, which will provide access to our own tests so you can fix any remaining problems. However, *that will be too late for you to test most of the compiler*, so it is important that you write your own tests for all the features of COOL covered in this MP.

The directory contains its own Makefile. Some of the targets it provides are:

- `make file.ll`: compile the Cool program `file.cl` to LLVM assembly.
- `make file.bc`: create an LLVM bytecode file from `file.ll`
- `make file.s`: create a machine code assembly file from `file.s`
- `make file.exe`: create a linked executable from `file.s` (and `coolrt.c`, for MP 3)
- `make file.out`: execute `file.exe` and put the output in `file.out`
- `make file.verify`: verify your LLVM code obeys LLVM language rules.
- `make file-opt.bc`: create an optimized LLVM bytecode file from `file.exe.bc`. This is just so you can see whether your code can be optimized effectively by available techniques in LLVM.
- `make clean`: delete all generated files.

To be sure that you generate correct LLVM code you should call the LLVM verification target with every program that you generate. You can do this by running “`make file.verify`” as described above. See the target `%.verify` in `mp2/Makefile.common` for the command used (or use “`make file.verify -n`”).

5 Designing the Code Generator

There are many possible ways to write the code generator. One reasonable strategy is to perform code generation in three passes; this is the strategy used by our solution and by the skeleton code. The first pass decides the object layout for each class, i.e. which LLVM data types to create for each class, and generates LLVM constants for all constants appearing in the program. Using this information, the second and third passes recursively walk each feature and generate the allocas and the LLVM code, respectively, for each expression.

There are a number of things you must keep in mind while designing your code generator:

- You should have a clear picture of the runtime semantics of Cool programs. The semantics are described informally in the first part of the *CoolAid*, and a precise description of how Cool programs should behave is given in Section 12 of the manual.

- You should have a clear picture of LLVM instructions, types, and declarations.
- Think carefully about how and where objects, let-variables, and temporaries (intermediate values of expressions) are allocated in memory. The next section discusses this issue in some detail.
- You should generate unoptimized LLVM code, using a simple tree-walk similar to the one we discussed in class. Focus on generating reasonably efficient local code for each tree node, e.g., wherever possible, avoid extra casts, use `getelementptr` to index into objects (i.e., to compute the address of a structure field), use appropriate aggregate types, etc.
- Ignore the garbage collection requirement of Cool. You don't have to implement it. Just insert calls to the function `i8* @malloc(i32)` to allocate heap objects whenever needed, and never free these objects.

In order to run a COOL program and inspect its result, your compiler should add a `main()` function to the generated LLVM module. This function should call `@Main_main()` and print the result. It should look like the following or equivalent:

```
define i32 @main() {
entry:
    %tpm.0 = call i32 @Main_main( )
    %tpm.1 = getelementptr [25 x i8], [25 x i8]* @.str, i32 0, i32 0
    %tpm.2 = call i32(i8*, ... ) @printf( i8* %tpm.1, i32 %tpm.0 )
    ret i32 0
}
```

The `.exe` target will fail until your compiler generates a valid assembly file that defines a `main` function with the right signature, and `.out` targets may fail rather than capture output if your generated assembly has sufficiently serious errors.

You should generate this function explicitly using LLVM IR features. To make this easier for you, we've provided a skeleton routine called `CgenClassTable::code_main()` in `cgen.cc`.

Your code generation phase executable `cgen` takes a `-c` flag to generate debugging information. This is set whenever you define `debug` true in your Makefile (the default). Using this flag merely causes `cgen_debug` (a global variable) to be set. Adding the actual code to produce useful debugging information is up to you. See the project *README* for details.

6 Representing Objects and Values in COOL

A major part of your compiler design is to develop the correct representation and memory allocation policies for objects and values in COOL, including explicit variables, heap objects, and temporaries. In MP 2, you only need to be concerned with `Int` and `Bool` values, and only as variables or temporaries (not heap objects).

Here are the guidelines you should follow:

- Values of primitive types should be represented directly as virtual registers (of types `i32` and `i1`) in your generated code, always for MP 2 and in most cases for MP 3.

- The only time an `Int` or `Bool` must live on the heap in your compiler is if an actual object operation needs to be performed on it. Ordinary arithmetic operations (`+`, `<`, etc.) are not object operations. Assignment of a value to an `Int` or `Bool` variable is not an object operation. There are no object operations in MP 2 (in MP 3, you will use boxing and unboxing for those operations on `Int` and `Bool`).
- A corollary is that the return value from `@Main_main` must be an `i32` and not `i32*`.
- Think of *let*-variables as pointers to values (objects): this is the correct interpretation for COOL because the same variable can be assigned different values (and so must point to different heap objects) at different places within its *let*-block. Since a *let*-variable has a local scope, we can allocate it in the current stack frame using the `alloca` instruction. Even if a *let*-variable is of a primitive type (`i32` or `i1`), we will just allocate it on the stack and let *mem2reg* promote it to an SSA register for us.

To enable this promotion, all `alloca` instructions should be placed in the *entry block* of the function. This placement ensures that your allocation is guaranteed to happen exactly once in each function invocation (even if the `let` is within a loop), improving the ability of optimizations to reason about it. We provide the declaration of a `make_alloca` function on expressions that could be implemented to perform the placement of `alloca` separately from the rest of the code-generation, to achieve the desired placement of `alloca`.

7 How to attack this project

To simplify your project, we strongly recommend you tackle it incrementally, taking the following steps in order to build your compiler. Make sure to test each portion of code as you complete it.

1. Start by generating the function `@main()` as described above, so that you can test your compiler even in the early stages of your work.
2. Start by implementing `Int` and `Bool` constants. At first, you can just generate them as the `i32` and `i1` LLVM primitives. *Test your compiler!*
3. Once you have constants, you can implement arithmetic and comparison operators. You can also implement block expressions at this time (e.g., `{1 + 2; 2 <= 1}`). *Test your compiler!*
4. Now try implementing `let`, following the allocation guidelines in the previous section. Use the environment to keep track of the binding from COOL variable names to memory locations (i.e., to LLVM `alloca/global/malloc` values). *You know what you need to do now!*
5. Next, implement assignment. Here, you will need to think about how the LHS and RHS are implemented, and what should be copied over.
6. Next, tackle `loop` and `if-then-else`. For these, you will need to learn more about LLVM `BasicBlocks`. In addition to the regular sources, the *Stacker* documentation contains some useful tips on using `BasicBlocks` to implement control structures.

The result of an `if-then-else` is a merge of the results of the two branches. You can allocate an `i32` or `i1` in the stack (depending on the type of the then-else branches) and then store a different result in each of the branches.

7. The final step: implement runtime error handling, if you haven't already. There are only a few cases you need to check, and they're listed in the back of the Cool manual.

For MP 2, the only possible error is divide-by-zero. Your program should call the function *abort()* if this happens. We've given you code to insert a declaration for *abort()* in the module.

8. Now test your compiler more thoroughly. You can use the Cool files we will give you in the *test-1* directory, but you should also make your own tests to stress individual cases.

8 What and how to hand in

You have to hand in all files that you modify in this MP. That is

- *cgen.cc*, *cgen.h*
- *cool-tree.handcode.h*
- *stringtab.handcode.h*
- *value_printer.cc*, *value_printer.h* (even if unchanged)
- *operand.cc*, *operand.h* (even if unchanged)

The four files, *value_printer.{cc,h}* and *operand.{cc,h}* should not need to be changed, but some students choose to modify them to support their code generator. Hand them in whether or not you change them.

Don't copy and modify any part of the support code! The provided files are the ones that will be used in the grading process.

The handin script for this project will be made available close to the deadline. As noted, we will give you enough time to test and fix any remaining bugs you missed during your own testing, but it will not be enough to fix major parts of your code generator.