

Programming Assignment V

Due Tuesday, April 18, 2006 at 11:59 PM
(Checkpoint due Tuesday, April 4, 2006 at 11:59 PM)

1 Introduction

In this assignment, you will implement a code generator for Cool. When successfully completed, you will have a fully functional Cool compiler!

The code generator makes use of the AST constructed in PA3 and static analysis performed in PA4. Your code generator should produce MIPS assembly code that faithfully implements *any* correct Cool program. There is no error recovery in code generation—all erroneous Cool programs have been detected by the front-end phases of the compiler.

As with the static analysis assignment, this assignment has considerable room for design decisions. Your program is correct if the code it generates works correctly; how you achieve that goal is up to you. We will suggest certain conventions that we believe will make your life easier, but you do not have to take our advice. As always, explain and justify your design decisions in the README file. This assignment is about twice the amount of the code of the previous programming assignment, though they share much of the same infrastructure. **Start early!**

Critical to getting a correct code generator is a thorough understanding of both the expected behavior of Cool constructs and the interface between the runtime system and the generated code. The expected behavior of Cool programs are defined by the operational semantics for Cool given in Section 13 of the *Cool Reference Manual*. Recall that this is only a specification of the meaning of the language constructs—not how to implement them. The interface between the runtime system and the generated code is given in *The Cool Runtime System*. See that document for a detailed discussion of the requirements of the runtime system on the generated code. There is a lot of information in this handout and the aforementioned documents, and you need to know most of it to write a correct code generator. *Please read thoroughly.*

You will build on your code generator that you develop here for PA6.

You may work in a group of one or two people. The `submit` program will ask you to specify group members when you turn in your assignment.

2 Files and Directories

As with previous assignments, we strongly suggest you use CVS to manage the project with your partner. Follow the instructions on the course website to set up a CVS repository. After you have set up your repository and checked-out a working directory for PA5, `cd` to that directory. If you are using C++, run this command to get the initial skeleton:

```
gmake -f ~cs164/assignments/PA5/Makefile source
```

For Java, type:

```
gmake -f ~cs164/assignments/PA5J/Makefile source
```

(notice the “J” in the path name).

As usual, there are other files used in the assignment that are symbolically linked to your directory or are included from `~cs164/include` and `~cs164/src`. You should not modify these files. Almost all of these files have been described in previous assignments.

We now describe the most important files for each version of the project.

2.1 C++ Version

This is a list of the files that you may want to modify. Most of the other files, you are probably familiar with from the previous assignments. See the `README` file for details about the additional files.

- `cgen.cc`

This file will contain almost all your code for the code generator. The entry point for your code generator is the `program_class::cgen(ostream&)` method, which is called on the root of your AST. Along with the usual constants, we have provided functions for emitting MIPS instructions, a skeleton for coding strings, integers, and booleans, and a skeleton of a class table (`CgenClassTable`). You can use the provided code or replace it with your own inheritance graph from PA4.
- `cgen.h`

This file is the header for the code generator. You may add anything you like to this file. It provides classes for implementing the inheritance graph. You may replace or modify them as you wish.
- `emit.h`

This file contains various code generation macros used in emitting MIPS instructions among other things. You may modify this file.
- `cool-tree.h`

As usual, these files contain the declarations of classes for AST nodes. You can add field or method declarations to the classes in `cool-tree.h`. The implementation of methods should be added to `cgen.cc`.
- `cgen_supp.cc`

This file contains general support code for the code generator. You will find a number of handy functions here. Add to the file as you see fit, but don't change anything that's already there.
- `example.cl`

This file should contain a test program of your own design. Test as many features of the code generator as you can.
- `README`

This file will contain the write-up for your assignment. It is critical that you explain design decisions, how your code is structured, and why you believe your design is a good one (i.e., why it leads to a correct and robust program). It is part of the assignment to explain things in text as well as to comment your code.
- `checkpoint`

You should answer the questions in this file, and submit these answers to us by the checkpoint due date. See Section 5.

2.2 Java Version

This is a list of the files that you may want to modify. Most of the other files, you are probably familiar with from the previous assignments. See the README file for details about the additional files.

- `CgenClassTable.java` and `CgenNode.java`

These files provide an implementation of the inheritance graph for the code generator. You will need to complete `CgenClassTable` in order to build your code generator. You can use the provided code or replace it with your own inheritance graph from PA4.

- `StringSymbol.java`, `IntSymbol.java`, and `BoolConst.java`

These files provide support for Cool constants. You will need to complete the method for generating constant definitions.

- `cool-tree.java`

This file contains the definitions for the AST nodes. You will need to add code generation routines (`code(PrintStream)`) for Cool expressions in this file. The code generator is invoked by calling method `cgen(PrintStream)` of class `program`. You may add new methods, but do not modify the existing declarations.

- `TreeConstants.java`

As before, this file defines some useful symbol constants. Feel free to add your own as you see fit.

- `CgenSupport.java`

This file contains general support code for the code generator. You will find a number of handy functions here including ones for emitting MIPS instructions. Add to the file as you see fit, but don't change anything that's already there.

- `example.cl`

This file should contain a test program of your own design. Test as many features of the code generator as you can.

- README

This file will contain the write-up for your assignment. It is critical that you explain design decisions, how your code is structured, and why you believe your design is a good one (i.e., why it leads to a correct and robust program). It is part of the assignment to explain things in text as well as to comment your code.

- checkpoint

You should answer the questions in this file, and submit these answers to us by the checkpoint due date. See Section 5.

3 Design

Before continuing, we suggest you read *The Cool Runtime System* to familiarize yourself with the requirements on your code generator imposed by the runtime system. The checkpoint will then help you make sure that you solidify your understanding of this.

In considering your design, at a high-level, your code generator will need to perform the following tasks:

1. Determine and emit code for global constants, such as prototype objects.
2. Determine and emit code for global tables, such as the `class_nameTab`, the `class_objTab`, and the dispatch tables.
3. Determine and emit code for initialization method for each class.
4. Determine and emit code for each method definition.

There are many possible ways to write the code generator. One reasonable strategy is to perform code generation in two passes. The first pass decides the object layout for each class, particularly the offset at which each attribute is stored in an object. Using this information, the second pass recursively walks each feature and generates stack machine code for each expression.

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

- Your code generator must work correctly with the Cool runtime system, which is explained in the *Cool Runtime System* manual.
- You should have a clear picture of the runtime semantics of Cool programs. The semantics are described informally in the first part of the *Cool Reference Manual*, and a precise description of how Cool programs should behave is given in Section 13 of the manual.
- You should understand the MIPS instruction set. An overview of MIPS operations is given in the `spim` documentation, which is in the course reader and on the class web page.
- You should decide what invariants your generated code will observe and expect (i.e., what registers will be saved, which might be overwritten, etc). You may also find it useful to refer to information on code generation in the lecture notes and portions of the text, primarily ASU Chapter 9.

You do *not* need to generate the same code as `coolc`. `Coolc` includes a very simple register allocator and other small changes that are not required for this assignment. The only requirement is to generate code that runs correctly with the runtime system.

3.1 Runtime Error Checking

The end of the Cool manual lists six errors that will terminate the program. Of these, your generated code should catch the first three—dispatch on void, case on void, and missing branch—and print a suitable error message before aborting. You may allow SPIM to catch division by zero. Catching the last two errors—substring out of range and heap overflow—is the responsibility of the runtime system in `trap.handler`. See Figure 4 of the *Cool Runtime System* manual for a listing of functions that display error messages for you.

3.2 Garbage Collection

To receive full credit for this assignment, your code generator must work correctly with the generational garbage collector in the Cool runtime system. The skeletons contain functions `code_select_gc` (C++) and `CgenClassTable.codeSelectGc` (Java) that generate code that sets GC options from command line flags. The command-line flags that affect garbage collection are `-g`, `-t`, and `-T`. Garbage collection is disabled by default; the flag `-g` enables it. When enabled, the garbage collector not only reclaims memory, but also verifies that “-1” separates all objects in the heap, thus checking that the program (or

the collector!) has not accidentally overwritten the end of an object. The `-t` and `-T` flags are used for additional testing. With `-t` the collector performs collections very frequently (on every allocation). The garbage collector does not directly use `-T`; in `coolc` the `-T` option causes extra code to be generated that performs more runtime validity checks. You are free to use (or not use) `-T` for whatever you wish.

For your implementation, the simplest way to start is not to use the collector at all (this is the default). When you decide to use the collector, be sure to carefully review the garbage collection interface described in the *Cool Runtime System* manual. Ensuring that your code generator correctly works with the garbage collector in *all* circumstances is not trivial.

4 Testing and Debugging

You will need a working scanner, parser, and semantic analyzer to test your code generator. You may use either your own components or the components from `coolc`. By default, the `coolc` components are used. To change that, replace the `lexer`, `parser`, and/or `semant` executable (which are symbolic links in your project directory) with your own scanner/parser. Even if you use your own components, it is wise to test your code generator with the `coolc` scanner, parser, and semantic analyzer at least once because we will grade your project using `coolc`'s version of the other phases.

You will run your code generator using `mycoolc`, a shell script that “glues” together the generator with the rest of compiler phases. Note that `mycoolc` takes a `-c` flag for debugging the code generator; using this flag merely causes `cgen_debug` (a global variable in the C++ version and a static field of class `Flags` in the Java version) to be set. Adding the actual code to produce useful debugging information is up to you. See the project README for details.

4.1 Coolaid

`Coolaid` is a tool to statically verify some basic correctness properties of the MIPS assembly code produced from Cool source. `Coolaid` will check that the assembly code is “well-typed” with respect to the Cool typing rules, just like the Java bytecode verifier checks that the bytecode output by a Java compiler is type safe. Aside from checking the safety of the output code, this tool can greatly benefit the development process of the code generation phase. Since the compiler front-end ensures that the source program is well-typed and that the type system guarantees certain safety properties, we expect that those properties should also hold on the resulting assembly code. If `Coolaid` is not able to verify some particular safety property, then a likely cause is a bug in the compiler itself.

We have tested `Coolaid` on 8200 programs generated by student compilers from the Spring 2002 and the Spring 2003 offering of CS164 at UC Berkeley. The standard testing procedure (which runs the generated code on the MIPS simulator `spim`) used for grading found errors in 2500 of those. Often the error messages were in the form of garbled output or output that did not match the expected output, whereas `Coolaid` was able to give precise error messages where the unsafe operation was performed that most likely caused the garbled output. `Coolaid` also finds errors in 1300 of the 5700 programs that pass the testing procedure! These were errors that the testing procedure missed because the sample input for the code did not exercise the bad code portions. In the Spring 2004, the students were given access to `Coolaid`. As a result, the mean autograder score of PA5 for the Spring 2004 was 79%, almost a letter grade improvement from the previous years (67% for the Spring 2002 and 69% for the Spring 2003)! See the *Coolaid Reference Manual* to get started using `Coolaid`. Note `Coolaid` imposes some additional requirements on the generated code to be able to check it; these requirements are also discussed in the *Cool Runtime System* manual.

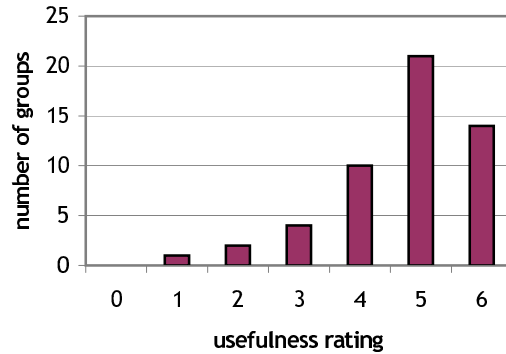


Figure 1: Student feedback from the 2004 class about the usefulness of Coolaid. 0 means “counterproductive” and 6 means “can’t imagine developing/debugging a compiler without it.”

Extra Credit. Coolaid is research technology that is being deployed in CS164. We would like to encourage the use of Coolaid both because we would like to receive feedback on our tool and because we truly believe it will be helpful in developing your compilers. As such, we will consider giving extra credit for bug reports on Coolaid. Whether and how much extra credit will be awarded will be decided on a case-by-case basis.

Feedback. If you use Coolaid, please help us by answering the following questions in the README:

1. Rate Coolaid’s usefulness on a scale from 0 to 6 where 0 means that it was counter-productive, 3 means that it neither helped nor hindered you, and 6 means that you can’t imagine developing/debugging a code generator without it.
2. What aspects of Coolaid particularly helped you?
3. What aspects of Coolaid would you like to see improved?
4. Describe bugs, if any, that Coolaid helped you find or fix?
5. Do you have any additional comments you would like to share with us?

4.2 Spim and XSpim

The executables `spim` and `xspim` are simulators for MIPS architecture on which you can run your generated code. The program `xspim` works like `spim` in that it lets you run MIPS assembly programs. However, it has many features that allow you to examine the virtual machine’s state, including the memory locations, registers, data segment, and code segment of the program. You can also set breakpoints and single step your program. The documentation for `spim/xspim` is in the course reader or on the course web page.

Warning. One thing that makes debugging with `spim` difficult is that `spim` is an interpreter for assembly code and not a true assembler. If your code or data definitions refer to undefined labels, the error shows up only if the executing code actually refers to such a label. Moreover, an error is reported only for undefined labels that appear in the code section of your program. If you have constant data definitions that refer to undefined labels, `spim` won’t tell you anything. It will just assume the value 0 for such undefined labels.

5 Checkpoint

As with PA4, this is a large project. The checkpoint file in your PA5 directory has several questions you should answer to make sure you understand how to interface with the runtime system. You should answer these questions in the file provided and submit them by running “`submit PA5-checkpoint`”. Your submission will be counted as a written assignment and be graded pass/fail.

6 Final Submission

Make sure to complete the following items before submitting to avoid any penalties.

- Include your write-up in README. Please include your feedback about Coolaid in the README.
- Include your test cases that test your code generator in `example.cl`.
- Make sure all your code for the code generator is in
 - `cool-tree.h`, `cgen.h`, `cgen.cc`, `cgen_supp.cc`, and `emit.h` for the C++ version; or
 - `cool-tree.java`, `CgenClassTable.java`, `CgenNode.java`, `CgenSupport.java`, `BoolConst.java`, `IntSymbol.java`, `StringSymbol.java`, `TreeConstants.java`, and additional `.java` files you may have added for the Java version.
- Be sure to answer ‘yes’ to the submission prompt for files that contain your code.