Programming Assignment IV[1]
A Semantic Analyzer for COOL

Due Friday, November 14, 2003

1. Introduction

In this assignment you will implement the static semantics of Cool. You will use the abstract syntax trees (AST) built by the parser to check that a program is in conformance with the Cool specification. Your static semantic component should reject erroneous programs; for correct programs, it must gather certain information for use by the code generator. The output of the semantic analyzer will be an attributed AST for use by the code generator.

This assignment has much more room for design decisions than previous assignments. Your program is correct if it checks programs against the specification. There is no "right" way to do the assignment, but there are wrong ways. There are a number of standard practices which we think make life easier and we will try to convey them to you. However, what you do is largely up to you. Whatever you decide to do, be prepared to justify and explain your solution.

You will need to refer to the typing rules, identifier scoping rules, and other restrictions of Cool as defined in the *CoolAid* manual. You will also need to add methods and data members to the AST class definitions for this phase. The functions the tree package provides are documented in the *Tour of Cool Support Code* (C++ version) and online (Java version):

http://www-inst.eecs.berkeley.edu/~cs164/public_html-sp03/javadoc

There is a lot of information in this handout, and you need to know most of it to write a working semantic analyzer. Please read the handout thoroughly.

2. Files and Directories

We strongly suggest you use CVS for this assignment. This assignment is larger than the previous ones, so version control may be especially helpful.

If you are using C++, run this command:

```
gmake -f ~icom4029/cool/assignments/PA4/Makefile source
```

For Java, type:

```
gmake -f ~icom4029/cool/assignments/PA4J/Makefile source
```

(notice the \J" in the path name). This will create several symbolic links in the correct directory. You should not edit the files pointed to by these links. In fact, if you make and modify private copies of these files, you may find it impossible to complete the assignment. See the instructions in the README file.

This is a list of the files that you will need to modify:

---

[1] This programming assignment is essentially the same one used at UC Berkeley CS 164 in the Spring of 2003. Used with permission.

- `semant.cc` (C++ version)

This file contains a start on a semantic analysis phase, written in C++. Put the code for your semantic analysis phase in this file. The skeleton only includes things required to correctly meet the interface with the code generator. (This interface is discussed in detail below.) Very little of real interest is included. The semantic analyzer is invoked by calling method `semant()` of class `program`.

- `semant.h` (C++ version)

This file is the header file for `semant.cc`.

- `cool-tree.handcode.h and cool-tree.h` (C++ version)

These files are where user-defined extensions to the abstract syntax tree nodes are placed. You can modify either file, but you will probably find it easiest to modify `cool-tree.h`. Do not modify the existing declarations.

- `cool-tree.java` (Java version)

This file contains the definitions for the AST nodes. You will need to add the code for your semantic analysis phase in this file. The semantic analyzer is invoked by calling method `semant()` of class `program`. Do not modify the existing declarations.

- `ClassTable.java` (Java version)

This class is a placeholder for some useful methods (including error reporting and initialization of basic classes). You may wish to enhance it for use in your analyzer.

- `TreeConstants.java` (Java version)

This file defines some useful symbol constants.

- `good.cl bad.cl`

These files test a few semantic features. You should add tests to ensure that `good.cl` exercises as many legal semantic combinations as possible and that `bad.cl` exercises as many kinds of semantic errors as possible. It is not possible to exercise all possible combinations in one file; you are only responsible for achieving reasonable coverage. Explain your tests in these files and put any overall comments in the README file.

- `README`

This file will contain the write-up for your assignment. For this assignment it is critical that you explain design decisions, how your code is structured, and why you believe that the 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. Inadequate README files will be penalized more heavily in this assignment, as the README is the major guideline we have to understanding your code.

- `checkpoint`

Submit your preliminary design writeup using this file. See Section 10.

As usual, there are other files used in the assignment that are symbolically linked to your directory or are included from ~icom4029/cool/include/PA4. You should not modify these files. Almost all of these files have been described in previous assignments.

3. Testing the Semantic Analyzer

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

You will run your semantic analyzer using `mysemant`, a shell script that "glues" together the analyzer with the parser and the scanner. Note that `mysemant` takes a -s flag for debugging the analyzer; using this flag merely causes `semant_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.

Once you are confident that your semantic analyzer is working, try running `mycoolc` to invoke your analyzer together with other compiler phases. You should test this compiler on both good and bad inputs to see if everything is working. Remember, bugs in the semantic analyzer may manifest themselves in the code generated or only when the compiled program is executed under `spim`.

4. Tree Traversal

As a result of assignment 3, your parser builds abstract syntax trees. The method dump with types, defined on most AST nodes, illustrates how to traverse the AST and gather information from it. This algorithmic style - a recursive traversal of a complex tree structure - is very important, because it is a very natural way to structure many computations on ASTs.

Your programming task for this assignment is to 1) traverse the tree, 2) manage various pieces of information that you glean from the tree, and 3) use that information to enforce the semantics of Cool. One traversal of the AST is called a "pass". You will probably need to make at least two passes over the AST to check everything.

You will most likely need to attach customized information to the AST nodes. To do so, you may edit `cool-tree.h` (C++) or `cool-tree.java` (Java) directly. In the C++ version, any method definitions you wish to add should go into `semant.cc`.

5. Inheritance

Inheritance relationships specify a directed graph of class dependencies. A typical requirement of most languages with inheritance is that the inheritance graph be acyclic. It is up to your semantic checker to enforce this requirement. One fairly easy way to do this is to construct a representation of the type graph and then check for cycles. A well known method for accomplishing this is to conduct a depth-first traversal of the graph.

In addition, Cool has restrictions on inheriting from the basic classes (see the manual). It is also an error if class A inherits from class B but class B is not defined.

The project skeleton includes appropriate definitions of all the basic classes. You will need to incorporate these classes into the inheritance hierarchy.

We suggest that you divide your semantic analysis phase into two smaller components. First, check that the inheritance graph is well-defined, meaning that all the restrictions on inheritance are satisfied. If the

inheritance graph is not well-defined, it is acceptable to abort compilation (after printing appropriate error messages, of course!). Second, check all the other semantic conditions. It is much easier to implement this second component if one knows the inheritance graph and that it is legal.

6. Naming and Scoping

A major portion of any semantic checker is the management of names. The specific problem is determining which declaration is in effect for each use of an identifier, especially when names can be reused. For example, if **i** is declared in two let expressions, one nested within the other, then wherever **i** is referenced the semantics of the language specify which declaration is in effect. It is the job of the semantic checker to keep track of which declaration a name refers to.

As discussed in class, a symbol table is a convenient data structure for managing names and scoping. You may use our implementation of symbol tables for your project. Our implementation provides methods for entering, exiting, and augmenting scopes as needed. You are also free to implement your own symbol table, of course.

Besides the identifier self, which is implicitly bound in every class, there are four ways that an object name can be introduced in Cool:

- attribute definitions
- formal parameters of methods
- let expressions
- branches of case statements

In addition to object names, there are also method names and class names. It is, of course, an error to use any name that has no matching declaration.

Remember that neither, classes, methods, nor attributes need be declared before use. Think about how this affects your analysis.

7. Type Checking

Type checking is another major function of the semantic analyzer. The semantic analyzer must check that valid types are declared where required. For example, the return types of methods must be declared. Using this information, the semantic analyzer must also verify that every expression has a valid type according to the type rules. The type rules are discussed in detail in the *CoolAid* and the course lecture notes.

One difficult issue is what to do if an expression doesn't have a valid type according to the rules. First, an error message should be printed with the line number and a description of what went wrong. It is relatively easy to give informative error messages in the semantic analysis phase, because it is generally obvious what the error is. We expect you to give informative error messages. Second, the semantic analyzer should attempt to recover and continue. A good semantic analyzer will avoid cascading errors using any of several standard techniques. We do expect your semantic analyzer to recover, but we do not expect it to avoid cascading errors. A simple recovery mechanism is to assign the type Object to any expression that cannot otherwise be given a type (we used this method in `coolc`).

8. Code Generator Interface

For the semantic analyzer to work correctly with the rest of the `coolc` compiler, some care must be taken to adhere to the interface with the code generator. We have deliberately adopted a very simple, naive interface to avoid cramping your creative impulses in semantic analysis. However, there is one thing you must do. For every expression node, its type field must be set to the Symbol naming the type inferred by your type checker. This Symbol must be the result of the `add_string` (C++) or `addString` (Java) method of the `idtable`. The special expression `no_expr` must be assigned the type `No_type` which is a predefined symbol in the project skeleton.

9. Output and Grading

For incorrect programs, the output of semantic analysis is error messages. You are expected to recover from all errors except for ill-formed class hierarchies. You are also expected to produce complete and informative errors. Assuming the inheritance hierarchy is well-formed, the semantic checker should catch and report all semantic errors in the program. When in doubt, use `coolc` as a guide in determining what informative error messages should say. Your error messages need not be identical to those of `coolc`.

We have supplied you with a simple error reporting method `ClassTable::semant_error()` (C++) and `ClassTable.semantError()` (Java). This routine takes a filename and the AST node where the error occurred, and it returns the error stream after it has printed an error header. The filename should be the file in which the error occurs. The parser ensures that `Class_` nodes store the file in which the class was defined (recall that class definitions cannot be split across files). In an error message, the line number of the error message is obtained from the AST node where the error is detected and the file name is obtained from the enclosing class.

For correct programs, the output is a type-annotated abstract syntax tree. You will be graded on whether your semantic phase correctly annotates ASTs with types and on whether your semantic phase works correctly with the coolc code generator.

You are also expected to program in good, structured style. You should spend some time thinking about the class definitions you will use. Grading will be out of 50 points, with 4 points for the README, 4 for the test cases in `good.cl` and `bad.cl`, 2 for comments, and 40 points for correctly typechecking our tests.

10. Checkpoint

The semantic analysis phase is by far the largest component of the compiler so far. Our solution is approximately 1300 lines of well-documented C++. You will find the assignment easier if you take some time to design the semantic checker prior to coding. Ask yourself:

- What requirements do I need to check?
- When do I need to check a requirement?
- When is the information needed to check a requirement generated?
- Where is the information I need to check a requirement?

By the checkpoint due date, you must submit a text file called checkpoint that contains your detailed answers to these questions. Submit it by running "submit PA4-checkpoint". Your submission will be counted as a written assignment and be graded pass/fail. As with other written assignments, no late checkpoints will be accepted.

If you can answer these questions for each aspect of Cool, implementing a solution should be straightforward.

At a high level, your semantic checker will have to perform the following major tasks:

1. Look at all classes and build an inheritance graph.
2. Check that the graph is well-formed.
3. For each class
    (a) Traverse the AST, gathering all visible declarations in a symbol table.
    (b) Check each expression for type correctness.

This list of tasks is not exhaustive; it is up to you to faithfully implement the specification in the manual.