

UNIVERSITY OF CALIFORNIA  
Department of Electrical Engineering  
and Computer Sciences  
Computer Science Division

CS 164  
Spring 2008

P. N. Hilfinger

**Project #2: Static Analyzer for Pyth (version of 5 March 2008)**

**Due:** Wednesday, 2 April 2008

The second project picks up where the last left off. Beginning with the AST we produced in Project #1, you are to perform a number of static checks on its correctness, annotate it with information about the meanings of identifiers, and perform some transformations to forms that would have been inconvenient to produce in the parser.

You may implement your solution in either C++ or Java. Your job is to hand in a program (the parser and its testing harness), including adequate internal documentation (comments), and a thorough set of test cases, which we will run both against your program and everybody else's.

Following a description of the annotations you must provide and some rewritings of the AST you must perform, we've accumulated various rules that are stated or implicit in the Pyth documentation, and that must all get checked before a program gets through with static semantic analysis. A program that advances to code generation will obey all these rules, as well as having the semantic annotations we describe.

## 1 Input and Output

We'll provide a glue script that takes input arguments, generates an AST out of the specified input plus a *standard prelude*, described below, and feeds this AST to your semantic analyzer. The output from your analyzer will be a revised tree with annotations that indicate which declaration applies to which identifier and also give the type of each subexpression.

We'll provide the tedious software for reading and outputting the tree. Since it is difficult to test the output of your project directly against any kind of standard (many different ASTs encode the same program after all), we'll augment `unparse` to show your annotations.

## 2 The Standard Prelude

The term *standard prelude* is used to describe the set of predefined things that a language provides. To vastly simplify your task in this project, we'll provide the standard prelude as a

file of Pyth definitions that get prepended to the source program (that's why we've had that provision for inputting multiple source files, in fact). It will define predefined types (like **Any** and **Dict**), methods (like `__add__`), and constants (like **None**). Since the parser will already have turned operators into method calls, you won't have to do anything special about them.

Indeed, you can *almost* treat the standard prelude as just part of the source program text. There are just a few unavoidable little glitches, in which the handling of the standard prelude differs from elsewhere:

- The built-in types will be defined with **class** declarations, which they normally can't be.
- The superclass of **Any** is **Any** (a class is not normally allowed to extend itself).
- The built-in types can use other built-in types as supertypes (normally, you have to use subtypes of **Object**).
- The type **Void** is a subtype of everything, which normal types cannot be.

For the purposes of this project, we'll differentiate these special-case types by the rule that all types up through and including **Object** are special. Types other than the special cases must obey the usual Pyth rules: they must have **Object** as a supertype and must reference only previously declared classes.

### 3 Output Format

The output ASTs differ from input ASTs in these respects:

- Most identifiers will have an extra annotation at the end:

```
(id F N "name" (decl D))
```

where  $D \geq 1$  is an integer *declaration index*.

- Nodes other than id nodes that denote expressions (things with values) will have a type annotation of the form **(stype T)** added to the end, where *T* is a type node such as produced in Project #2 for type declarations, with annotations. For example,

```
(not 0 10 (id 0 10 "x" (decl 42))
  (stype (type (id "Bool" (decl 5)))))
```

The expression operators are

```
and call compare dict float getrhs int list not or select string tuple
```

plus a few new ones introduced below.

- The **program** node will contain a third child consisting of a list of all declaration nodes (in order by index).

- The analyzer will apply several re-writes to the tree so as to produce a tree that would have been inconvenient for the parser to produce.

There is one declaration index (and corresponding declaration node) for each distinct declaration in the program: each class definition (including built-in types), local variable, parameter, constant definition, method definition, instance variable, and class variable. Table 1 shows the formats of the declaration nodes.

## 4 Rewritings

Because the parser lacks the necessary information (or finds it inconvenient to do the right thing), the semantic analyzer must do some re-writing of the trees that come from the parser, in addition to the annotations we've already described. Table 2 shows rewrites that happen after annotation.

### 4.1 Various Restrictions

The context-free grammar from Project #1 allows certain things that Pyth disallows. Some of these *could* be enforced in the grammar, but weren't required to be for convenience.

1. A **break** or **continue** statement may occur only in a loop (**for** or **while**).
2. A **return** statement may only occur in a **def**.
3. An identifier, *C*, that refers to a class (in the scope of a **class** definition) may be used only in the following ways:
  - a. In the inheritance clause of another class definition;
  - b. On the left of an attribute reference: *C.x*;
  - c. As the function name in a call (i.e., in an object creation): *C()*.
  - d. As a type name in a type assertion.

Thus, constructs like `x = C`, `f (C)`, or `if C` are all illegal.

4. An inheritance clause in a class must reference a class completely defined previously in the program.
5. A class def (static method) must appear immediately within the body of a class (and therefore not inside any method definition).

If you choose to use your solution to Project #1 as the parser and lexer in Project #2, then any of these that it already handles need not be checked. In any case, all of these errors must be reported in or before the static semantic analyzer.

**Table 1:** Declaration nodes. The analyzer adds a third child to the **program** node consisting of a list of the declaration nodes for a program in order by index. In each case,  $N$  is the declaration index, unique to each declaration node instance. Types are represented as in Project #1.

Node	Meaning
<code>(localdecl N "I" P T)</code>	Local variable named $I$ of type $T$ . $P$ is the declaration index of the enclosing function, if any, or 0 for variables at the outer level.
<code>(paramdecl N "I" P T)</code>	Parameter named $I$ of type $T$ defined in a function whose declaration index is $P$ .
<code>(constdecl N "I" P T)</code>	A constant named $I$ of type $T$ and defined inside the class or function whose declaration index is $P$ (0 for constants at the outer level).
<code>(instancedecl N "I" P T)</code>	Instance variable named $I$ of type $T$ defined in the class with declaration index $P$ .
<code>(funcdecl N "I" P T [<math>f_1 \dots f_n</math>] [<math>v_1 \dots v_m</math>])</code>	A class method or ordinary function named $I$ of type $T$ , defined in class or function with declaration index $P$ (0 for the outer level), whose formal parameters have declaration indices $f_1, \dots, f_n$ and that contains local declarations (of local variables, nested functions, constants) whose declaration indices are $v_1, \dots, v_m$ .
<code>(methoddecl N "I" P T [<math>f_1 \dots f_n</math>] [<math>v_1 \dots v_m</math>])</code>	An instance method. The arguments are the same as for <code>funcdecl</code> .
<code>(classdecl N "I" P [<math>m_1 \dots m_n</math>])</code>	Class declaration for class named $I$ . $P$ is the declaration index of the parent type. The $m_i$ are the declaration numbers of the class members that are introduced or overridden in the class.
<code>(nativeclassdecl N "I" P [<math>m_1 \dots m_n</math>])</code>	Built-in class declaration. Same as class declaration, but indicates that this is one of the built-in classes from the standard prelude.

**Table 2:** Rewriting Rules. Annotations and source-position information are elided for brevity. New nodes should be given the source positions of the nodes they replace.

Original Node	Rewrite
<code>(call (id i) [E<sub>1</sub> ... E<sub>n</sub>])</code> <i>where i denotes an instance method on T<sub>1</sub>, the static type of E<sub>1</sub>.</i>	<code>(callmethod (id i) [E<sub>1</sub> ... E<sub>n</sub>])</code>
<code>(call (select (id i) E<sub>0</sub>) [E<sub>1</sub> ... E<sub>n</sub>])</code> <i>where i denotes an instance method and E<sub>0</sub> denotes a value (this does not apply if E<sub>0</sub> denotes a type, because of the next rewrit- ing rule).</i>	<code>(callmethod (id i) [E<sub>0</sub> E<sub>1</sub> ... E<sub>n</sub>])</code>
<code>(select (id i) T)</code> <i>where either T denotes a type (as opposed to a value), or i denotes a class method or constant.</i>  <code>(setselect (id i) T v)</code> <i>where T denotes a type (as opposed to a value).</i>	<code>(id i)</code>     <code>(assign (id i) v)</code>
<code>E<sub>0</sub></code> <i>where E<sub>0</sub> appears in a context that requires it to have a subtype of T and the static type of E<sub>0</sub>, while feasible for T, is not a subtype of T.</i>	<code>(coerce E<sub>0</sub> (stype T))</code>

## 4.2 Scope rules

In Pyth, the declarative regions (called *namespaces* in Python documentation) are as follows:

- The *global region*, consisting of the source file that is input to the compiler and the built-in definitions or *standard prelude* (in Python, these are two nested regions);
- A *class region* for each class definition;
- A *local region* including the parameters and body of each function or method definition (**def**).

Local regions nest inside each other and inside the global region. Class regions nest in the global region.

The scope of a **def** is the entire declarative region in which it occurs, except where hidden in a nested scope. Likewise, the scope of a class definition is the entire global region. A name defined by **def** may not be redefined in the same declarative region, and classes may not be redefined at all. Since an assignment in Pyth, as in Python, causes a declaration, this that any two of the following are illegal together:

```
f = 3
def f (x): ... # Illegal redefinition.
def f (y): ... # Another illegal redefinition
class f: ...   # And still another
```

The scope of a parameter is the entire function definition in which it occurs, except where hidden by inner declarations of the same name.

If there is an assignment or augmented assignment to a variable in a declarative region, or if it is implicitly assigned by being used as a control variable in a **for** statement, then there is (one) implicit definition of that variable in the innermost region containing the assignment, *unless* there is a **global** declaration of that variable in the region. The scope of such an implicit definition is the entire innermost region containing it (before and after the assignment), except where hidden by inner declarations, as usual. There is at most one implicit definition of a variable generated by this rule; multiple assignments have the same effect as one. If there is a **global** definition of a name in a certain region, then all uses of that name in the region (again, except where hidden) refer to the definition of the name in the global region. A name defined as global must also be defined by **def** or (implicitly) by assignment in the global region. Thus, the following program is illegal:

```
def f ():
    global a    # Illegal: no assignment to a in the global region
    a = 3
# a = 0        # Program WOULD be legal if this statement were uncommented.
f ()
print a        # Illegal: a is not defined.
```

The following are illegal:

```

class foo:
    global a      # Illegal: no assignment or def a in the global region
    global b
    def b (): ... # Illegal: b already declared as global.
def b (): ...
# def a (): ...  # global a WOULD be legal if this were uncommented

```

The nesting rules already given imply that if a name is *not* assigned to or **def**ed in some function body, then any use of it in that body refers to a **def**, or to an implicit definition caused by an assignment, in some surrounding region; there must be one for the program to be legal.

A variable is implicitly defined by the presence of an assignment statement even if, during execution, that assignment never happens. Thus, the following is a statically *legal* program that may cause a runtime error when executed (if **a** is used but never set):

```

def f (x):
    if x > 3:
        a = "Answer is %d" % x
    print a[0]    # Fails if a is not set (and so has value None).

```

In other words, your static analyzer never considers the particular values a variable will take, only their types.

When the innermost declarative region surrounding an assignment is a class definition, the assigned-to variable is an *instance variable* of the class. The class also inherits all instance variables of its parent (if any). An assignment to an instance variable of the parent does not create a second instance variable of the same name; it refers to the parent's variable. These are the only legal instance variables (in Python, you can introduce new instance variables into a class or class object by assignment to an attribute at any time). In addition, the assigned-to variable is also defined as a *class variable* (static in Java). Thus,

```

class Stuff:
    x = 13
    def g (self, y): self.x = y

z = Stuff ()
z: Stuff
z.g (42)
print Stuff.x  # Prints 13
print z.x      # Prints 42

```

When your static checker sees **Stuff.x**, it will know that **Stuff** is the name of a class. It should also be able to determine whether **x** is defined as a class variable of that class. The same is true when you see **z.x** or **z.g**. However, were the type declaration **z: Stuff** missing, your semantic analyzer should report an error, since the static type of **z** would then be **Any**, and there would be no static guarantee that fields **x** and **g** existed.

A class may **def** (not **class def**) a function that has a **def** in its parent, which has the effect of overriding the parent's definition.

## 5 Types

By default, any defined entity in Pyth has type **Any**. A **defed** instance method defined in class  $T$  has (by default) a type  $(T, \text{Any}, \dots) \rightarrow \text{Any}$ , where the number of types to the left of the arrow is the number of parameters declared for the function. Other **defed** functions have by default type  $(\text{Any}, \dots) \rightarrow \text{Any}$ . A Pyth type assertion ascribes a type to a defined name. The name must be declared somewhere in the same innermost declarative region that contains its type assertion, and successive assertions for the same name must be compatible. When you ascribe a type to a **defed** function, the number of parameters must be the same. The types of the parameters are ascribed to the formal parameters. Formal parameters may not be given types in the body of the function or method.

The job of your analysis is to determine the most specific type you can for each expression and declaration in the program, and to find any uses of a name that *must* be illegal. For this purpose, you only need to use information gleaned from type assertions, **defs**, and the language rules about certain basic constructs of the language: string literals are of type **String**, integer literals of type **Int**, tuples of type **Tuple**, list displays of type **List**, and dictionary displays of type **Dict**. If you know the type of a function, you can determine the type of a call to the function. Since the AST translates many Pyth constructs into calls, this fact will take care of most type checking for you.

As an example, your analyzer should be able catch these illegalities:

```
def f (x, y): x + y
print f (3)           # ERROR: number of arguments doesn't match f's type
class Foo:
    x: Int
    x = 3
S: String
S = "a"
S = 3                 # ERROR: wrong type
Foo.x = S             # ERROR: wrong type
```

Your program is *not* required to catch this:

```
f = 3
f(2)                  # There is no type assertion or def for f, so you don't
                      # know whether it's a function.
3[1]                  # __getitem__ is defined for all types.
a: List
a = [ "x", "y", "z" ]
a[0] + 1              # Only know that a is a List, not what's in it
a[5]                  # Don't understand about index bounds.
```

The specific legality rules we want checked are:

1. In a call  $f(E_1, \dots, E_n)$ , either  $f$  must have type **Any** (i.e., unknown) or it must have a function type with  $n$  arguments.



2. Also, if  $f$  has the type  $(T_1, \dots, T_n) \rightarrow T_0$ , then the known static type of each  $E_i$  must be *feasible* for type  $T_i$ . We say that type  $A$  is feasible for type  $F$  if either  $A$  is a subtype of  $F$  or  $F$  is a subtype of  $A$ . Unlike Java, this allows many programs where the compiler cannot tell for certain that a particular call or assignment is allowed. What it does not allow is passing something known to be an `Int` as a `String` parameter.
3. In an assignment  $x = E$ , the type of  $E$  must be feasible for that of  $x$ .
4. In an `assigns` node (produced for assignments where the left-hand operator is a tuple) the type of  $E$  must be `Any`, `Tuple`, `List`, or `Dict`.
5. In an attribute reference  $X.y$ , the attribute  $y$  is known to be an attribute of  $X$  (an instance variable for a class, e.g.).
6. No assignment is allowed to a `defed` entity.

Most other interesting cases are handled by the function-call rule.

## 6 Output and Testing

The output of the program is again a textual representation of the AST (plus error messages). It's even more difficult than in the first project to check that your program's output is correct. Error tests will be of particular importance: you must make sure that breaking any of the rules causes an error. We will augment `unparse` to handle annotated ASTs and to reconstruct the source program with little annotations indicating which definition connects to each use, and giving type information for defined quantities. Once again, testing will be an important part of your grade.

## 7 What to turn in

The directory you turn in (under the name `proj2-n` in your `tags` directory) should contain a file `Makefile` that is set up so that

```
gmake
```

(the default target) compiles your program and

```
gmake check
```

runs all your tests against your program. We'll put sample Makefiles (for C++ and Java) in the `~cs164/hw/proj2` directory; feel free to modify at will as long as these two commands continue to work.

## 8 What We Supply

We'll supply a framework to read and write ASTs and declarations (which, of course you may modify however you wish). We'll also maintain `unparse`.

## 9 Assorted Advice

What, you haven't started yet? First, review the Pyth language, and start writing test cases. You get points for thorough testing and documentation, and it should not be difficult to get them, so don't put this off to the last minute!

Again, be sure to ask us for advice rather than spend your own time getting frustrated over an impasse. By now, you should have your partner's phone number at least. Keep in regular contact.

Be sure you understand what we provide. The skeleton classes actually do quite a bit for you. Make sure you don't reinvent the wheel.

Do not feel obliged to cram all the checks that are called for here into one method! Keep separate checks in separate methods. To the extent possible, introduce and test them one at a time.

Keep your program neat at all times. Keep the formatting of your code correct at all times, and when you remove code, remove it; don't just comment it out. It's much easier to debug a readable program. Afraid that if you chop out code, you'll lose it and not be able to go back? That's what Subversion is for. Archive each new version when you get it to compile (or whenever you take a break, for that matter). This will allow you to go back to earlier versions at will.

Write comments for classes and functions before you write bodies, if only to clarify your intent in your own mind. Keep comments up to date with changes. Remember that the idea is that one should be able to figure how to use a function from its comment, without needing to look at its body.

You *still* haven't started?