

Thesis for the degree of Master of Science

An executable operational semantics for Python

Gideon Joachim Smeding

January, 2009
INF/SCR-08-29



Universiteit Utrecht

Center for Software Technology
Dept. of Information and Computing Sciences
Universiteit Utrecht
Utrecht, The Netherlands

Supervisors:
dr. Andres Löh
prof. dr. S.D. Swierstra

Abstract

Programming languages are often specified only in an informal manner; in the available documentation, the language behavior is described by examples and text. Only the implementation, a compiler or interpreter, describes the exact semantics of constructs.

Python is no different. It is described by an informal manual and a number of implementations. No systematic, formal descriptions of its semantics are available.

We developed a formal semantics for a comprehensive subset of Python called *minpy*. The semantics are described in literate Haskell. The source files are compiled to an interpreter as well as the formal specifications in this document. In a sense, this document is the documented source code of the interpreter.

Contents

1	Introduction	7
1.1	Python	7
1.2	Formal semantics	8
1.3	Contributions	10
1.4	Overview	10
2	Preliminaries	11
2.1	Lists	11
2.2	Mappings	11
3	Abstract syntax	13
3.1	Expressions	13
3.2	Statements	14
4	Transforming an abstract machine	15
5	Theory of objects and classes	17
5.1	The method resolution order	17
5.2	Mappings of inherited attributes	20
5.3	Classes as types	21
6	Basic rewrites rules	23
6.1	Blocks	23
6.2	Pass and expression statements	24
7	Variables	25
7.1	Assignment statements	25
7.2	Variable expressions	26
7.3	Delete statements	26
8	Functions and generators	27
8.1	Defining functions and generators	27
8.2	Executing functions	29
8.3	Executing generators	33
9	Classes and objects	37
9.1	Class definition	37
9.2	Object creation	39
9.3	Attribute access	41
9.4	Attribute assignment	42
9.5	Attribute deletion	43

Contents

10	Exceptions	45
10.1	The raise statement	45
10.2	The try-except statement	47
10.3	The try-finally statement	48
11	Control structures	51
11.1	The if statement	51
11.2	The while statement	52
11.3	The for statement	54
11.4	The break statement	55
11.5	The continue statement	56
12	Print statements	57
13	Operators	59
13.1	Unary operators	59
13.2	Binary operators	60
13.3	Container operators	62
14	Dynamic code execution	65
14.1	The exec statement	65
14.2	The import statement	66
15	Conclusion	69
15.1	Literate programming	69
15.2	The interpreter	70
15.3	The test suite	72
15.4	Future work	73
A	Builtin values	75
A.1	The base objects <code>object</code> and <code>type</code>	75
A.2	The function and generator	76
A.3	Integer values	76
A.4	Booleans	78
A.5	Strings	78
A.6	Lists	80
A.7	Tuples	82
A.8	The tuple class	82
A.9	Dictionaries	83

1 Introduction

This introduction is organized as follows. First we review Python with respect to formal semantics, and define a scope for our work. Then we discuss executable operational semantics and our approach to it. Finally, we list the main contributions, and give an overview of the rest of this thesis.

1.1 Python

Python is an imperative, dynamic, object-oriented programming language originally developed by Guido van Rossum at the CWI in the Netherlands in the 1980s. Since the first publication of the code in 1991, Python has gained much popularity and now is a major programming platform.

A formal semantics for Python

There are no formal semantics for Python. CPython is the de facto reference implementation. Although it maintains high coding standards, CPython is not written with legibility as its primary focus.

The PyPy [11] implementation of Python is written in Python itself. To be exact, it is written in a restricted form of Python called RPython [1]. The introduced restrictions facilitate static analyses such as typing, and enable good run-time efficiency.

Because it is easy to read for Python programmers, it has been suggested that PyPy might one day become the reference implementation of Python. However, even for the restricted form of Python, no systematic documentation of the operational semantics exists.

Despite the lack of a formal semantics, Python appears to be a particularly good candidate for formal specification, for a number of reasons.

- Python is a relatively simple language. It has a limited number of constructs, most of which are found in many other languages. Python introduces no radically new features, but has a unique combination of features found in its predecessors.
- Despite being a simple language, it is a widely used language which is still gaining in popularity. As a language and platform, Python has proven to be mature: many sizable projects use it to great success.
- The language is quite well documented. There is an extensive reference manual [15] and all past changes to the language have been documented in Python enhancement proposals (PEPs) [14].
- There are a number of independent, mature implementations of Python available [3, 6, 7, 11]. Semantics can be used to compare interpreters and maintain language conformism.

These semantics target Python version 2.5. The behaviour of the CPython implementation, not the documentation, is taken as the definition of Python.

Scope of the semantics

Of course we would have liked to document all of Python, but time constraints did not allow that. Therefore the scope of these semantics is limited to a subset of Python that we call *minpy*. The language *minpy* includes virtually all syntactical constructs of Python, but leaves out much of the syntactic sugar (see Chapter 3 for an overview of the abstract syntax).

While most of the concepts introduced by the syntax are included in these semantics, much of the language's features have been ignored:

- The standard library has not been specified wherever possible. Some elements of the library are needed to support other features. For example, exception handling naturally requires the `Exception` class.
- Garbage collection is not modeled by these semantics. Apart from managing the storage of objects in memory, garbage collection affects the operational semantics by calling an object's finalizer, which can be specified by a programmer. Typically different implementations have their own garbage collectors and some even allow a user to control it.
- Python's support for multi-threading is ignored. While multi-threading is an interesting subject by itself, it is simply beyond the scope of this project. It could also be argued that threads are not truly part of the language, since they are only supported through the interpreter's libraries.
- The ability of Python to interact with the 'outside world' is not described. In practice, the 'outside world' is represented by other libraries. Interaction between two different languages is achieved through what is sometimes called a foreign function interface.
- The specification of simple types in this document, such as integers and strings, do not describe the exact behaviour of CPython (see Chapter A). The semantics of simple types tend to differ in the small details. Types and their associated operators usually have the semantics of the platform underpinning the interpreter. For example, Jython uses Java's integers while CPython's integers are based on C's integer types.
- These semantics only work with the so called new style classes that have been introduced in CPython version 2.2 [4]. The old-style classes had been kept for backwards compatibility only, and will not be in future versions of Python.
- Python has various reflexive features. For example, it is possible to inspect the dictionary that implements objects, and the body of a function can be replaced at runtime. These features are not covered by our semantics.

1.2 Formal semantics

Formal semantics seek to precisely describe the meaning of a programming language using mathematical constructs. All aspects of programming languages can be formally specified, but usually a formal semantics describes the run-time behaviour of programs.

Operational semantics

A number of different styles of formal semantics exist. Operational semantics as introduced by Plotkin [10], describe language semantics in terms of a state transition system. Using an abstract machine for the state, the operational semantics will closely resemble an interpreter. Because a state machine uses well known constructs like stacks and heaps, the semantics will be comparatively easy to understand.

In general, formal semantics improve our understanding of a language to its finest details, simplifying reasoning about programs at an abstract level. Operational semantics are of significant practical value as well:

- The detailed documentation provided by an operational semantics facilitate the creation of interpreters of compilers. Especially if accompanied by a test suite, standardization is much simpler to achieve and maintain.
- Program analyses can be designed in a more systematic fashion and formal reasoning can be used to verify properties of analyses. More generally, many language tools are easier to develop with an operational semantics at hand.
- A language can be extended and altered more safely, since interaction of extensions with the original language can be analyzed systematically. For example backwards compatibility can be guarded more closely.

While the case for operational semantics is easy to make, some disadvantages of formal semantics in general also apply to operational semantics, albeit to a lesser degree.

- Formal semantics are often perceived to be an ivory tower matter, i.e., be of little value to neither the users nor the developers of a language. The following anecdote illustrates this nicely. When inquiring after previous work on formal semantics on the Python mailing list, someone answered “[...] I don’t think Python culture operates like that very much.”
- Formal semantics are often defined separately from the actual implementation. Whether or not a compiler or interpreter actually adheres to the defined semantics is unclear and difficult to judge. Very few compilers or interpreters have been proven correct with respect to the semantics.
- It has been argued that formal semantics might inhibit evolution of a language [5]; the maintenance costs of formal semantics could discourage experimentation and extension.

An executable operational semantics

The listed shortcomings of formal semantics can be eliminated by making the operational semantics executable. In other words, by writing the semantics in such a way, that it can be compiled into an interpreter.

An executable operational semantics would be of immediate practical value as an interpreter. Furthermore it simplifies maintenance of the operational semantics, as the semantics and implementation are developed simultaneously. An executable operational semantics enables experimentation and simplifies language evolution.

1 Introduction

The operational semantics in this document have been written in the programming language Haskell, or more specifically literate Haskell, a variant that allows Haskell code to be mixed with \LaTeX . The Haskell code has been formatted using `lhs2TeX`[9] and some simple scripts. Thus the sources used to produce this document, can also be compiled into an interpreter for *minpy*.

The source code has been reformatted in such a way, that no knowledge of Haskell is required. One might even forget that the semantics were written in a programming language at all, as the rules look like common mathematical equations.

Normally writing the semantics of a complete language is a tedious and error prone process. Writing an executable operational semantics on the other hand is exciting, as one's progress is clearly reflected in the interpreter. The strong type system of Haskell prevents many kinds of mistakes, and we can compare the behaviour of our executable semantics to that of Python with test cases. While testing cannot prove the correctness of the semantics, it does justify confidence in the semantics' quality.

1.3 Contributions

The main contributions of this master's thesis are the following:

1. Firstly, we have described the operational semantics of a significant subset of Python;
2. secondly, we have created an interpreter from the same sources as the operational semantics;
3. and finally, we have created a test suite that compares the behaviour as described by our operational semantics to the CPython implementation, or other implementations.

1.4 Overview

This thesis is organised as follows. First, we set the stage for the semantics by introducing some notational conventions in Chapter 2, describing the abstract syntax of *minpy* in Chapter 3, and introducing the abstract machine model in Chapter 4. Then, in Chapter 5 we introduce the object model of *minpy*.

The remaining chapters, except for the conclusion, contain the semantic rules that describe the behaviour of specific statements and expressions. We start with some basic rules in Chapter 6, followed by the rules describing variables in Chapter 7, functions and generators in Chapter 8, classes and objects in Chapter 9, exception handling in Chapter 10, control structures in Chapter 11, printing in Chapter 12, operators in Chapter 13, and finally the `exec` and `import` statements in Chapter 14.

Finally, in Chapter 15, we discuss the results of this work and the lessons that we learned in the process.

2 Preliminaries

Before discussing the semantics themselves we introduce the notational conventions used for lists and mappings.

2.1 Lists

In these semantics we use many lists. Lists are sequences which are ordered collections of elements. For example, a list $[3, 5, 2, 5]$ contains the elements 3, 5, 2, and 5 in that order. Lists can also be empty, which would be denoted as pair of empty brackets.

By convention we overline names of lists, e.g. \bar{a} would be a list of addresses. Individual elements of \bar{a} are referred to by their index starting at 1. For example, q_2 refers to the number 4 in the list $\bar{q} = [1, 4, 3]$. The length of a list $|\bar{q}|$ is defined as the number of elements in the list. The operator $:$ prepends a list with a new element and the operator $\#$ concatenates two lists.

2.2 Mappings

A mapping is a partial function that maps keys to values. For example, arrays and hash tables are mappings. Python itself has native support for mappings, which it calls dictionaries.

We use a number of special operators and notations to describe, alter and inspect maps. These operators have the same semantics for all mappings, regardless of their specific contents.

The list of key-value pairs $[k_1 \mapsto v_1, \dots, k_n \mapsto v_n]$ represents a mapping m of keys $k_{1..n}$ to values $v_{1..n}$ respectively. Each key $k \in m$ maps to a value $v = m(k)$. Thus, the key-value pairs in the mapping are unique. An empty mapping is denoted as \emptyset , which is never used for empty sets.

The operator \oplus combines two mappings A and B to create a new mapping that consists of the key-value pairs of both A and B . For duplicate keys, the key-value pairs of the left hand side of \oplus take precedence over the right hand side. Formally defining this operator, we have:

$$(A \oplus B)(k) = \begin{cases} B(k) & \text{if } k \in B \\ A(k) & \text{otherwise} \end{cases}$$

The operator \ominus removes a key-value pair, identified by its key, from a mapping. If the key is not in the domain of the mapping, nothing changes. Formally defining this operator, we have:

$$(A \ominus k)(j) = \begin{cases} A(j) & \text{if } k \neq j \\ \perp & \text{otherwise} \end{cases}$$

2 Preliminaries

3 Abstract syntax

This chapter describes the abstract syntax of *minpy*. First we will introduce the syntax of expressions and operators, followed by statements and blocks. The semantics of each syntactical construct will be described in the following chapters.

The syntax definitions in this chapter are reformatted Haskell data definitions that mimic concrete Python syntax. Lists of a syntax element A are denoted as $\langle A \rangle^*$.

Because the concrete syntax of Python has already been specified by the reference manual [13], it will not be discussed here. The parser used by the *minpy* interpreter only implements a minimal subset of Python's concrete syntax.

The abstract syntax clearly shows some of the limitations of *minpy* compared to full Python. For example, list comprehension expressions are missing and try-except statements are limited to a single except clause. To keep the specification simple, we have made various optional elements mandatory. For example, the else branch in the if-then-else statement is not optional in this abstract syntax.

3.1 Expressions

The expression syntax specification indicates no binding preferences. Any ambiguity will be resolved by parenthesis. For example $3 + 2 * 5$ will have to be written as $3 + (2 * 5)$.

```
Expr ::= Name                -- Variable , see Rule 7.3
        | Expr( $\langle \textit{Expr} \rangle^*$ ) -- Function call , see Rule 8.2
        | Expr.Name          -- Attribute access , see Rule 9.10
        | Expr[Expr]         -- Slice access , see Rule 13.11
        | Expr BinOp Expr   -- Binary operator , see Rule 13.5
        | UnaryOp Expr       -- Unary operator , see Rule 13.1
        | yield Expr         -- Yield expression , see Rule 8.13
        | Int                -- Literal integer , see Rule A.5
        | Bool               -- Literal boolean , see Rule A.10
        | String             -- Literal string , see Rule A.13
        | [ $\langle \textit{Expr} \rangle^*$ ]   -- Literal list , see Rule A.21
        | ( $\langle \textit{Expr} \rangle^*$ )   -- Literal tuple , see Rule A.30
        | { $\langle \textit{Expr} : \textit{Expr} \rangle^*$ } -- Literal dictionary, see Rule A.35
```

```
BinOp ::= + | - | * | / | % | ** | // | == | != | < | <= | > | >= | is | in | and | or
```

```
UnaryOp ::= not | -
```

3.2 Statements

The command line interface of Python accepts single statements. Full programs as well as modules, consist of a single block of statements –a list of statements– in a text file.

We often append a semicolon to statements to distinguish statements, and especially expression statements, from expressions. The concrete syntax of Python supports this use of semicolons as well, but does not require it.

Some statements, so called compound statements, contain blocks. In the concrete syntax these blocks are delimited by their indentation level (see Python’s reference documentation). In this abstract syntax we delimit blocks using curly brackets where ambiguity arises. Specifically the empty block is denoted as an empty pair of curly brackets.

<i>Stmt ::= Expr</i>	-- expression statement, see Rule 6.5
<i>Name = Expr</i>	-- variable assignment , see Rule 7.1
<i>Expr.Name = Expr</i>	-- attribute assignment, see Rule 9.15
<i>Expr[Expr] = Expr</i>	-- slice assignment , see Rule 13.14
<i>del Name</i>	-- variable deletion , see Rule 7.4
<i>del Expr[Expr]</i>	-- slice deletion , see Rule 13.18
<i>del Expr.Name</i>	-- attribute deletion , see Rule 9.19
<i>def Name((Name)*): Block</i>	-- function definition , see Rule 8.1
<i>return Expr</i>	-- return , see Rule 8.8
<i>print(Expr)</i>	-- print , see Rule 12.1
<i>if Expr : Block else : Block</i>	-- if-then-else , see Rule 11.1
<i>while Expr : Block</i>	-- while , see Rule 11.4
<i>for Name in Expr : Block</i>	-- for , see Rule 11.8
<i>class Name(Expr) : Block</i>	-- class definition , see Rule 9.1
<i>try : Block except Expr , Name : Block</i>	-- try-catch , see Rule 10.6
<i>try : Block finally : Block</i>	-- try-finally , see Rule 10.10
<i>raise Expr</i>	-- raise , see Rule 10.1
<i>import (Name)*</i>	-- import , see Rule 14.6
<i>pass</i>	-- pass , see Rule 6.4
<i>break</i>	-- break , see Rule 11.14
<i>continue</i>	-- continue , see Rule 11.15
<i>exec Expr in Expr</i>	-- exec , see Rule 14.1

Blocks in functions, classes and modules also define a variable scope, the semantics of which are discussed in chapters 7, 8, 9, and 14. There are some restrictions on the syntax of scoping blocks that are not expressed by the abstract syntax definitions. It could be argued that these are not syntactical restrictions, but Python raises a `SyntaxError` for programs that fail the restrictions:

- The return and yield statements are limited to the scope of a function. In other words, they can only occur in function definitions or in a while, for, try, or if statement inside the function.
- A function scope can contain either a return or a yield statement, but not both.
- The continue and break statements can only occur in for or while loops, or in a nested if or try statement.

4 Transforming an abstract machine

The operational semantics are defined by state transitions of an abstract machine. A machine state $\langle \Theta, \Gamma, S, \iota \rangle$ consist of a heap Θ , an environment stack Γ , a control stack S , and an instruction ι . State transitions are defined by rewrite rules that transform the machine state. These rewrite rules have the following shape.

$$\langle \Theta, \Gamma, S, \iota \rangle \Rightarrow \langle \Theta', \Gamma', S', \iota' \rangle$$

As the semantics of Python are deterministic, there is one, and only one, rule for each machine state. The states are discerned by pattern matches in the rules. In some cases however, pattern matches overlap. In these cases, the most specific pattern takes precedence over the others.

Heap

The heap is a mapping of addresses to values. Addresses are represented by natural numbers. Values include integers, strings, functions and objects.

For example, a heap $\Theta = [a_1 \rightarrow \text{"spam"}, a_2 \rightarrow 3, a_3 \rightarrow [a_1, a_2]]$, contains a string, a number, and a list at the addresses a_1 , a_2 , and a_3 respectively. A new heap $\Theta' = \Theta \oplus [a_1 \rightarrow \text{"eggs"}, a_4 \rightarrow 42]$ is defined as an update of the heap Θ : the string at a_1 is redefined and the value 42 is added at address a_4 .

Values are never removed from the stack. Because of this, the operational semantics presented here will never define a practical interpreter. In a full implementation of Python, unused values will be removed from the heap by a garbage collector.

Environment

The environment is a stack of addresses. The addresses point to environment mappings on the heap, each of which contains the bound variables in a single scoping block. Environment mappings are values that implement a mapping of variable names -represented by strings- to addresses.

The environment mappings on an environment stack can be merged to create the environment mapping Σ_Γ^Θ , that contains all variable bindings of the environment mappings on the environment stack. The merged environment mapping Σ_Γ^Θ is defined as:

$$\begin{aligned} \Sigma_\square^\Theta &= \emptyset \\ \Sigma_{\Gamma|\gamma}^\Theta &= \Sigma_\Gamma^\Theta \oplus \Theta(\gamma) \end{aligned}$$

For example, the environment Γ with a heap Θ binds the variables x and y to the values 1 and 2 respectively. The variable x is bound in both γ_1 and γ_2 , but the binding in γ_2 shadows the binding in γ_1 , i.e., $\Sigma_\Gamma^\Theta(x) = a_2$.

4 Transforming an abstract machine

$$\begin{aligned} \Gamma &= \square | \gamma_1 | \gamma_2 \\ \Theta &= [a_1 \rightarrow 1 \\ &\quad , a_2 \rightarrow 2 \\ &\quad , \gamma_1 \rightarrow ["x" \rightarrow a_2, "y" \rightarrow a_2] \\ &\quad , \gamma_2 \rightarrow ["x" \rightarrow a_1] \\ &\quad] \end{aligned}$$

Control stack

The control stack is a stack of continuation frames, i.e., frames that indicate ‘what to do next’. Many frames use a placeholder \circ , to indicate what part of an instruction is currently being executed or evaluated.

For example, the stack $\square | x = \circ$ consists of a single frame $x = \circ$ on top of the empty stack \square . The circle in the frame $x = \circ$ replaces the assignment’s right hand side to indicate that it is being evaluated. Once the right hand side expression has been evaluated, the assignment is executed.

Instruction

The instruction is a kind of program counter for the virtual machine. There are three different kinds of instruction: The instruction can be the expression, statement, or block that is being executed; it can be the result of the previously executed instruction; or an unwind instruction.

Side effects

Some rules have an effect on a hidden state not part of the abstract machine. Specifically I/O operations, such as printing to the screen, are typical side effects.

Because the outside world is not modelled in these semantics, side effects are informally described over the arrow of the rewrite rule. For example, a side effect q of a rule would be denoted as follows.

$$\langle \Theta, \Gamma, S, \iota \rangle \xrightarrow{q} \langle \Theta', \Gamma', S', \iota' \rangle$$

The side effect q occurs when the machine state is rewritten. Rules are never selected based on a side effect, i.e., the pattern match of a rewrite rule can not be influenced by the side effect of the rewrite rule.

The import statement (see Section 14.2) has no side effects, even though reading files usually has side effects. We assume however, that the imported files do not change during the execution of a program.

Source of a rewrite rule

The rewrite rules are written in Haskell and reformatted to yield formulas of the above form. For example, the source of Rule 6.3 is as follows:

```
rewrite      (State heap envs ( stack :|: BlockFrame b ) ( BwResult a )) = case b of
  EmptyBlock -> (state heap envs ( stack                ) ( BwResult a ))
  otherwise   -> (state heap envs ( stack                ) ( FwBlock   b ))
```


5 Theory of objects and classes

Python is an object-oriented language. Values, including built-in values such as integers, strings, and functions, are represented by objects. Objects are mappings of names to addresses. The addresses point to the object's data and the associated operations (functions).

Every object has exactly one class. Even classes, being objects themselves have a class. The class-of relations between object and their classes can be represented by a tree. At the root of such a tree is the class `type`, which is its own class. See for example Figure 5.1.

Every class has one or more base classes, with the exception of the object `object`, which has no bases itself. The bases of a class are sorted by their local precedence order. The bases of a class A , the bases of the bases, and so on, are the superclasses of A . The object `object` is a superclass of all classes.

The base-of relation between types can be represented by a directed (from class to base), acyclic graph with edges ordered by their local precedence order. In this inheritance graph, all paths eventually lead to the base object `object`. See for example Figure 5.2.

Both the class and bases are object attributes: the class of an object is stored in the attribute `__class__`, and the `__bases__` attribute contains a tuple of a class' base classes. The order of bases in the `__bases__` tuple define the local precedence order. Classes are differentiated from other objects, only by their `__bases__` attribute, which 'normal' objects lack.

All values $\Theta(a)$ on a heap Θ have an object representation Ω_a^Θ . The object representation is the mapping of names to values, i.e. the object's attributes. The value of an object is determined by the class of an object. For example, objects of class `type` or `object` have a value `None`, and objects of class `int` have a primitive integer value.

The object representations of built-in types and functions, are not stored on the heap. Built-ins have an immutable object representation $[_class_ \rightarrow a_C]$, where C is the class of the built-in type. Only a module value returns an object representation extended with its own mapping.

5.1 The method resolution order

A class inherits the attributes of its superclasses. Objects have access to the attributes of their class, including the inherited attributes, as if they are their own attributes. However, when classes have different definitions for the same attributes, it is unclear which takes precedence.

To resolve conflicts between inherited attributes, we define a so called method resolution order (MRO). The order of classes in the MRO determines which attribute overrides the other attributes of the same name. Despite the name, the MRO determines not only the resolution of methods (function attributes), but the resolution of all attributes.




Figure 5.1: Example tree showing class-of relations. The classes `object`, `A`, `B`, and `C` have class `type`, which is its own class. The objects `a` and `b` have classes `A` and `B` respectively.




Figure 5.2: Example inheritance graph. The classes `type`, `A`, and `B` inherit from `object`. The class `C` inherits `A` and `B`, in that order.




Figure 5.3: A (partial) extended inheritance graph. The marked arrow is added to the inheritance graph in Figure 5.2. Thus the class C has the MRO C, A, B, object .

The C3 algorithm

Python uses the C3 algorithm originally developed for Dylan [12, 2] to find an MRO that satisfies a number of requirements. Most importantly,

- The MRO observes local precedence order. For example, a class A that precedes a class B in the local precedence order of a class C , must precede B in the MRO of C .
- The MRO is monotonic. For example, a class A in the MRO of a class B , must precede B in any MRO that contains A .

To compute the C3 linearization, the inheritance graph is extended with edges for the local precedence order of bases: for each class with ordered list of bases \bar{C} we insert an edge from C_1 to C_2 , an edge from C_2 to C_3 , and so forth. The ordering of edges is no longer relevant in the extended inheritance graph.

The C3 algorithm produces a depth first topological sort of the extended inheritance graph, provided that the extended graph is acyclic. See for example Figure 5.3.

$$\begin{aligned} \text{MRO } \Theta(a) &= a : \bar{b} \sqcup \text{MRO } \Theta(b_1) \sqcup \text{MRO } \Theta(b_2) \sqcup \dots \\ \text{where } \bar{b} &= \Theta(\Omega_a^\Theta(\text{--bases--})) \end{aligned} \quad (5.1)$$

The MRO is defined using a left associative merge operator $\bar{a} \sqcup \bar{b}$ that merges the two topological sorts \bar{a} and \bar{b} . Note that, if \bar{a} contains a class that is also in \bar{b} , that class will occur only once in $\bar{a} \sqcup \bar{b}$.

$$\begin{aligned} [] \sqcup \bar{b} &= \bar{b} \\ \bar{a} \sqcup [] &= \bar{a} \\ (a : \bar{a}) \sqcup (b : \bar{b}) \text{ if } a \equiv b &= a : \bar{a} \sqcup \bar{b} \\ (a : \bar{a}) \sqcup (b : \bar{b}) \text{ if } a \not\equiv b \wedge a \notin \bar{b} &= a : \bar{a} \sqcup (b : \bar{b}) \\ (a : \bar{a}) \sqcup (b : \bar{b}) \text{ if } a \not\equiv b \wedge b \notin \bar{a} &= b : (a : \bar{a}) \sqcup \bar{b} \end{aligned} \quad (5.2)$$

Illegal inheritance graphs

Not all inheritance graphs are considered legal, e.g., cyclic extended inheritance graphs are not allowed. This sometimes has unintuitive consequences, see for example Figure 5.4.




Figure 5.4: The inheritance graph on the right has an illegal ordering of its base classes: the extended inheritance graph has a cycle $\text{object} \rightarrow B \rightarrow \text{object}$. The inheritance graph on the right is legal, because A inherits B before object and thus has no cycle in its extended inheritance graph.

Furthermore, all classes must have at least one base, with the exception of `object`, and classes cannot have duplicate bases. Only classes can serve as bases. Finally, a class can only have one built-in superclass (e.g. `int` or `list`).

5.2 Mappings of inherited attributes

As mentioned before, the MRO determines what attributes a class inherits. In order to lookup the attributes of an object Ω_a^\ominus , we define two derived mappings that encapsulate the inheritance of attributes.

A mapping of class attributes

We define the derived mapping of class attributes Φ_a^\ominus to contain the inherited attributes of the class of a and its own attributes. It is defined using a capital version of the join operator \oplus . The capital join operator $\bigoplus_\ominus \bar{a}$ merges the objects $\Omega_{a_1}^\ominus$ through $\Omega_{a_n}^\ominus$, where $n = |\bar{a}|$.

$$\begin{aligned} \bigoplus_\ominus [] &= \emptyset \\ \bigoplus_\ominus a : \bar{a} &= \bigoplus_\ominus \bar{a} \oplus \Omega_a^\ominus \end{aligned}$$

The derived mapping Φ_a^\ominus is defined as the joined mappings of the MRO of the class of a .

$$\begin{aligned} \Phi_a^\ominus &= \bigoplus_\ominus \text{MRO } \Theta(a_c) \\ &\text{where } a_c = \Omega_a^\ominus(_ _ \text{class} _ _) \end{aligned} \tag{5.3}$$

A mapping of all attributes

Finally, the full collection of attributes Υ_a^\ominus of an object Ω_a^\ominus is defined as the combination of its inherited attributes and the object's own attributes.

$$\Upsilon_a^\ominus = \Phi_a^\ominus \oplus \Omega_a^\ominus \tag{5.4}$$

5.3 Classes as types

In the object-oriented world the terms “type” and “classes” are often used interchangeably. In Python, the term “type” is somewhat ambiguously defined. We define two typing relations based on the class-of and base-of relations.

The subclass relation

Classes are said to be subclasses, or subtypes, of their base classes and the base classes thereof. We define the subclass relation $a_a <:\Theta a_b$ to hold if and only if a_a is a class, and a_b is in the MRO of a_a . If a class a_a is a subclass of a_b , we also say that a_b is a superclass of a_a .

$$a_a <:\Theta a_b \equiv \text{--bases--} \in \Omega_{a_a}^\Theta \wedge a_b \in \text{MRO } \Theta(a_a) \quad (5.5)$$

For example, the class C in Figure 5.3 is a subclass of itself, A , B , and `object`.

The instance-of relation

Before, we described the class-of relation as a relation between object and class: every object has exactly one class. The instance-of relation extends this notion to the superclasses of an object, i.e. an object is an instance of its class and all the superclasses of its class. We define the instance-of relation $a_o :_\Theta a_c$ to hold if and only if the class of a_o is a subclass of a_c .

$$a_o :_\Theta a_c \equiv \Omega_{a_o}^\Theta(\text{--class--}) <:\Theta a_c \quad (5.6)$$

For example, the object c in Figure 5.1 is an instance of C , A , B and `object` (assuming the inheritance graph of Figure 5.2).

5 *Theory of objects and classes*

6 Basic rewrites rules

6.1 Blocks

Blocks are sequences of statements. Python programs are represented by blocks. Compound statements, such as the while statement and module definitions, contain blocks.

Executing a non-empty block statement

A block is executed statement by statement: the first statement s of the block $s; b$ is executed, while the remaining statements in block b are put onto the stack for later execution.

$$\begin{aligned} & \langle \Theta, \Gamma, S, s; b \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S|b, s; \rangle \end{aligned} \tag{6.1}$$

Executing an empty block

An empty block is simply popped from the stack, returning the value `None`.

$$\begin{aligned} & \langle \Theta, \Gamma, S, \{\} \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S, a_{\text{None}} \rangle \end{aligned} \tag{6.2}$$

Rather than creating a new object on the heap, the rule returns a predefined address pointing to the built-in value `None`.

Continuing execution of a block

When the execution the previous statement has completed with result a , the block b on the stack is executed if it is not empty. If b is empty, the block frame is popped and we return a . Thus the result of a block's final statement is the result of the complete block.

$$\begin{aligned} & \langle \Theta, \Gamma, S|b, a \rangle \\ \Rightarrow & \begin{cases} \langle \Theta, \Gamma, S, a \rangle & \text{if } b = \{\} \\ \langle \Theta, \Gamma, S, b \rangle & \text{otherwise} \end{cases} \end{aligned} \tag{6.3}$$

6.2 Pass and expression statements

Pass statements

The pass statement affects neither the heap, the environment nor the stack and returns `None`.

$$\begin{aligned} & \langle \Theta, \Gamma, S, \text{pass}; \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S, a_{\text{None}} \rangle \end{aligned} \tag{6.4}$$

Expression statements

Expressions also serve as statements. Syntactically these expression statements are indistinguishable from normal expressions, therefore they are annotated with a suffix `;`

$$\begin{aligned} & \langle \Theta, \Gamma, S, e; \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S, e \rangle \end{aligned} \tag{6.5}$$

7 Variables

Variables are references to values on the heap. Variables can be assigned, evaluated (dereferenced), and deleted. Because variables are references, assignments and deletions change only the variable binding, not the values themselves. Multiple values can also refer to the same value. The program in Listing 7.1 shows how a references can be used to share values.

```
x = [1,2] # assigning a variable with a list
y = x    # aliasing the variable
x.append(3) # append 3 to the list
print(y)  # prints [1,2,3]
x = []    # reassigning x to the empty list
print(y)  # prints [1,2,3]
z = y
del y     # removes the binding for y
print(z)  # prints [1,2,3]
```

Listing 7.1: Assignments change the variable but not the value

7.1 Assignment statements

Executing an assignment statement

Executing the assignment $x = e$ starts with the evaluation of e . A continuation frame $x = \circ$, holding the variable name x , is placed on the stack.

$$\begin{aligned} & \langle \Theta, \Gamma, S, \quad, x = e; \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S | x = \circ, e \quad \rangle \end{aligned} \quad (7.1)$$

The circle \circ in the stack frame $x = \circ$ serves as a placeholder for the expression e that is being evaluated.

Binding a variable

When the evaluation of e is completed, the resulting address a is stored in the topmost environment mapping. Any previous binding of x in the mapping $\Theta(\gamma_1)$ will be overwritten.

$$\begin{aligned} & \langle \Theta \quad, \Gamma | \gamma_1, S | x = \circ, a \quad \rangle \\ \Rightarrow & \langle \Theta \oplus [\gamma_1 \rightarrow \Theta(\gamma_1) \oplus [x \rightarrow a]], \Gamma | \gamma_1, S \quad, a_{\text{None}} \rangle \end{aligned} \quad (7.2)$$

7.2 Variable expressions

Evaluating a variable expression

Evaluating a variable in general amounts to a simple lookup in the environment: if the variable is bound in the environment, we return it, otherwise an exception is raised.

Local variables (see Section 8.6), have slightly different semantics. A local variable can only be retrieved from the local environment, which is topmost in the environment stack.

Thus, the evaluation of a variable has four distinct cases:

1. If x is a local variable and is bound in the local environment $\Theta(\gamma_1)$, we lookup and return the address of x .
2. An `UnboundLocalError` exception is raised if x is a local variable but is unbound in the local environment.
3. If x not a local variable and is bound in the full environment $\Sigma_{\Gamma|\gamma_1}^\Theta$, we lookup and return the address of x .
4. Otherwise the variable is unbound and not local, so we raise a `NameError`.

$$\Rightarrow \left\langle \Theta, \Gamma|\gamma_1, S, x \right\rangle \Rightarrow \begin{cases} \left\langle \Theta, \Gamma|\gamma_1, S, \Theta(\gamma_1)(x) \right\rangle & \text{if } x \in \Theta(\gamma_1)_{loc} \wedge x \in \Theta(\gamma_1) \\ \left\langle \Theta, \Gamma|\gamma_1, S, \text{raise } a_{\text{UnboundLocalError}}() \right\rangle & \text{if } x \in \Theta(\gamma_1)_{loc} \wedge x \notin \Theta(\gamma_1) \\ \left\langle \Theta, \Gamma|\gamma_1, S, \Sigma_{\Gamma|\gamma_1}^\Theta(x) \right\rangle & \text{if } x \notin \Theta(\gamma_1)_{loc} \wedge x \in \Sigma_{\Gamma|\gamma_1}^\Theta \\ \left\langle \Theta, \Gamma|\gamma_1, S, \text{raise } a_{\text{NameError}}() \right\rangle & \text{otherwise} \end{cases} \quad (7.3)$$

7.3 Delete statements

Executing a variable deletion

The delete statement `del x` removes the variable x from the local environment γ_1 . As for the variable evaluation, local variables have a special case:

1. If x is bound in the local environment, we delete it.
2. An `UnboundLocalError` is raised if x is a local variable and not (yet) bound in the local environment.
3. Otherwise the variable must be unbound and not local, so we raise a `NameError`.

$$\Rightarrow \left\langle \Theta, \Gamma|\gamma_1, S, \text{del } x; \right\rangle \Rightarrow \begin{cases} \left\langle \Theta', \Gamma|\gamma_1, S, a_{\text{None}} \right\rangle & \text{if } x \in \Theta(\gamma_1) \\ \left\langle \Theta, \Gamma|\gamma_1, S, \text{raise } a_{\text{UnboundLocalError}}() \right\rangle & \text{if } x \notin \Theta(\gamma_1) \wedge x \in \Theta(\gamma_1)_{loc} \\ \left\langle \Theta, \Gamma|\gamma_1, S, \text{raise } a_{\text{NameError}}() \right\rangle & \text{otherwise} \end{cases} \quad (7.4)$$

where $\Theta' = \Theta \oplus [\gamma_1 \rightarrow \Theta(\gamma_1) \ominus x]$

8 Functions and generators

This chapter discusses functions and generators. Following the lifetime of a function, we define the semantics of functions in three steps: first the function definition, followed by the function call, execution, and the return. Next we discuss the operation of generators: (re)starting a generator, and 'returning' from a generator with the yield statement.

Functions in Python behave very similar to functions in other imperative language. A function takes a number of arguments, executes its body, and returns some value. For example Listing 8.1 shows how a factorial function is defined and called.

```
def fac(x) :
    if x < 1 :
        return 1
    else :
        return x * fac(x-1)

fac(5) # returns 120
```

Listing 8.1: A recursive factorial function

Generators are a unique concept of Python. They are defined using a normal function definition that contains one or more yield statements and no return statements (see Chapter 3). Such a generator function returns a generator when called.

When a generator's primitive function `__next__` is called, the generator executes the body of the generator function until a yield expression `yield e` is evaluated. Evaluating the yield expression causes the generator to suspend execution and return `e`. When `__next__` is called again, execution continues at the point where it left off, until either another yield expression is encountered or the end of the body is reached. At the end of its body, a generator stops and raises an exception.

For example Listing 8.2 shows how a generator is used to create a generator that returns the elements of a list one by one.

8.1 Defining functions and generators

A function definition `def $x_\lambda(\bar{x}) : b_\lambda$` consists of: the function's name x_λ , the list of parameter names \bar{x} , and the body b_λ . When executed, a new function object is added to the heap, the function name is bound in the environment, and `None` is returned.

Functions which body contains a yield expression, produce generators when called (see Section 8.3). The definition of these generator functions however, is no different from normal functions.

```

def createListGenerator(list) :
    try :
        i = 0
        while True :
            yield list[i]
            i = i + 1
    except IndexError, e :      # catch index out of bounds
        pass

it = createListGenerator([1,2])
it.next() # returns 1
it.next() # returns 2
it.next() # raises a StopIteration exception

```

Listing 8.2: A generator function that produces a list iterator

Executing a function definition statement

The new function object stored at a_λ consists of two parts: the map of attributes $[_class_ \rightarrow a_{\text{function}}]$, which only contains the class attribute, and a function closure value. The closure $\lambda \bar{x}. b_\lambda \dashv \Gamma_\lambda$ consists of: the function arguments \bar{x} , its body b_λ , and its environment Γ_λ .

$$\begin{aligned}
 & \langle \Theta, \Gamma | \gamma_1, S, \text{def } x_\lambda(\bar{x}) : b_\lambda; \rangle \\
 \Rightarrow & \langle \Theta', \Gamma | \gamma_1, S, a_{\text{None}} \rangle \\
 \text{where } & a_\lambda = \text{new address} \notin \Theta \\
 & \Gamma_\lambda = \text{if } \Theta(\gamma_1) \in \Omega \text{ then } \Gamma \text{ else } \Gamma | \gamma_1 \\
 & \Theta' = \Theta \oplus [a_\lambda \rightarrow ([_class_ \rightarrow a_{\text{function}}], \lambda \bar{x}. b_\lambda \dashv \Gamma_\lambda) \\
 & \quad , \gamma_1 \rightarrow \Theta(\gamma_1) \oplus [x_\lambda \rightarrow a_\lambda] \\
 & \quad]
 \end{aligned} \tag{8.1}$$

Class functions, recognized by the topmost environment (see Section 9.1), cannot access other class members directly. Therefore, the environment Γ_λ of class functions does not include the local environment γ_1 of the function definition statement. See Example 8.3.

```

x = 1
class A(object) :
    x = x*2 # define class variable A.x

    def foo(self) :
        return x

print(A().foo()) # prints 1
print(A.x)      # prints 2

```

Listing 8.3: Class variables are not in the scope of class variables

8.2 Executing functions

Evaluating calls

A call expression $e_f(\bar{e})$ executes a function e_f with arguments \bar{e} . Before the function is executed, e_f and \bar{e} are evaluated from left to right.

Evaluating the function expression

First, the function expression e_f is evaluated. A frame $\circ(\bar{e})$ is pushed onto the stack containing the remaining argument expressions.

$$\begin{aligned} & \langle \Theta, \Gamma, S, e_f(\bar{e}) \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S | \circ(\bar{e}), e_f \rangle \end{aligned} \quad (8.2)$$

Evaluating the first argument expression

Then, when e_f has been evaluated to a_f , we continue to evaluate the first argument e'_1 and push the remaining arguments \bar{e}' and a_f on the stack. If there are no arguments, a_f is executed (see Section 8.5).

$$\begin{aligned} & \langle \Theta, \Gamma, S | \circ(\bar{e}), a_f \rangle \\ \Rightarrow & \begin{cases} \langle \Theta, \Gamma, S | a_f(\circ, \bar{e}'), e \rangle & \text{if } \bar{e} = (e : \bar{e}') \\ \langle \Theta, \Gamma, S, a_f() \rangle & \text{otherwise} \end{cases} \end{aligned} \quad (8.3)$$

Evaluating the remaining argument expressions

Finally, the remaining arguments \bar{e} on the stack are evaluated one by one. When all argument expressions have been evaluated, the function is executed.

$$\begin{aligned} & \langle \Theta, \Gamma, S | a_f(\bar{a}, \circ, \bar{e}), a \rangle \\ \Rightarrow & \begin{cases} \langle \Theta, \Gamma, S | a_f(\bar{a} \# [a], \circ, \bar{e}'), e \rangle & \text{if } \bar{e} = (e : \bar{e}') \\ \langle \Theta, \Gamma, S, a_f(\bar{a} \# [a]) \rangle & \text{otherwise} \end{cases} \end{aligned} \quad (8.4)$$

Type checking an annotated function

The instance-of relation is not very informative for functions, since all functions are instances of the class `function`. However, we can describe a function more precisely by the types of arguments it accepts.

Function annotations describe the type of a function's arguments, which are checked right before execution of the function. For example the primitive addition function (see Rule A.6) is annotated (see Section A.3) to accept only instances of `int` as arguments.

Annotations are only introduced by rewrite rules and object definitions in these semantics. There is no annotation in the abstract syntax.

Type checking function arguments

An annotated function $a_f : \bar{c}$ consists of a function address a_f and a list of classes \bar{c} . We execute the annotated function call with arguments \bar{a} , if all arguments \bar{a} are instances of classes \bar{c} . There may be fewer classes in \bar{c} than there are arguments, but not more. If the arguments fail to satisfy the type, we raise an exception (see Chapter 10).

$$\begin{aligned} & \langle \Theta, \Gamma, S, a_f : \bar{c}(\bar{a}) \rangle \\ \Rightarrow & \begin{cases} \langle \Theta, \Gamma, S, a_f(\bar{a}) \rangle & \text{if } |\bar{a}| \geq |\bar{c}| \wedge \forall_{0 < i < |\bar{c}|} \bar{a}_i :_{\Theta} \bar{c}_i \\ \langle \Theta, \Gamma, S, \text{raise } a_{\text{TypeError}}() \rangle & \text{otherwise} \end{cases} \end{aligned} \quad (8.5)$$

Function closure execution

The execution of a fully evaluated call $a_\lambda(\bar{a})$ has four (successful) cases: a_λ is a function closure, a_λ is a generator function, a_λ is a callable object, or a_λ is a primitive function. The latter case is not handled in this chapter. Instead, every primitive function has a specific rewrite rule. See for example the primitive addition function which is handled in Rule A.6.

Execution of a function

The execution of a function call has four cases:

1. If a_f is an instance of `function` then the value of a_f is the closure $\lambda \bar{x}. b_\lambda \dashv \Gamma_\lambda$. When the number of arguments \bar{a} equals the number of parameters \bar{x} , and a_f is not a generator function, we execute the function body b_λ with the environment $\Gamma_\lambda | \gamma_1$. The current environment Γ is pushed on the stack in the return marker $\Gamma \vdash \circ_\lambda$.
2. If a_f is a closure with the correct number of arguments, and a_f is a generator function, we return a generator $\langle b_\lambda, \Gamma_\lambda | \gamma_1 \rangle$ (see Section 8.3). The generator's environment is the same as a function's environment would be, and the only element on its stack is the function body.
3. If a_f is not a closure, but an object with a class attribute `__call__`, we forward the call to that attribute. The function is instantiated with the object a_f .
4. Otherwise, the call had an incorrect number of arguments or a_f was neither a function nor a callable object, so we raise a `TypeError`.

$$\begin{aligned}
& \langle \Theta, \Gamma, S, a_f(\bar{a}) \rangle \\
\Rightarrow & \begin{cases} \langle \Theta', \Gamma_\lambda | \gamma_1, S | \Gamma \vdash \circ_\lambda, b_\lambda \rangle & \text{if } a_f :_\Theta a_{\text{function}} \wedge |\bar{a}| \equiv |\bar{x}| \wedge \text{yield} \notin b_\lambda \\ \langle \Theta'', \Gamma, S, a_{\text{gen}} \rangle & \text{if } a_f :_\Theta a_{\text{function}} \wedge |\bar{a}| \equiv |\bar{x}| \wedge \text{yield} \in b_\lambda \\ \langle \Theta, \Gamma, S, a_f \cdot a_{\text{call}}(\bar{a}) \rangle & \text{if } \neg (a_f :_\Theta a_{\text{function}}) \wedge \text{--call--} \in \Phi_{a_f}^\Theta \\ \langle \Theta, \Gamma, S, \text{raise } a_e() \rangle & \text{otherwise} \end{cases} \\
\text{where } & \gamma_1 = \text{new address} \notin \Theta \\
& \Theta' = \Theta \oplus [\gamma_1 \rightarrow ([\bar{x}_1 \rightarrow \bar{a}_1, \dots, \bar{x}_n \rightarrow \bar{a}_n], \text{lv}(b_\lambda) \cup \bar{x})] \\
& a_{\text{gen}} = \text{new address} \notin \Theta' \\
& \Theta'' = \Theta' \oplus [a_{\text{gen}} \rightarrow \langle b_\lambda, \Gamma_\lambda | \gamma_1 \rangle] \\
& a_{\text{call}} = \Phi_{a_f}^\Theta(\text{--call--}) \\
& a_e = a_{\text{TypeError}} \\
& \lambda \bar{x}. b_\lambda \dashv \Gamma_\lambda = \Theta(a_f)
\end{aligned} \tag{8.6}$$

The function's local environment γ_1 has two parts: the parameter-argument bindings $[\bar{x}_1 \rightarrow \bar{a}_1, \dots, \bar{x}_n \rightarrow \bar{a}_n]$ and the set of local variables $\text{lv}(b_\lambda) \cup \bar{x}$.

The local variables of a block are the variables used as a target in body of the function, i.e., the variable that can be bound in the local environment when the body is executed. Variable expressions and delete statements have different semantics for local variables (see Chapter 7).

Local variables

The set of local variables $\text{lv}(b)$ consists of all the variables that can be (un)bound during the execution a block b . Apart from the assign and delete statements, the for and try statements can also bind variables. There is no expression that binds variables. In that respect they are truly stateless, in contrast to statements.

$$\begin{aligned}
\text{lv}(\{\}) &= \emptyset \\
\text{lv}(x = e; & b) = x \cup \text{lv}(b) \\
\text{lv}(\text{del } x; & b) = x \cup \text{lv}(b) \\
\text{lv}(\text{while } e : b_x; & b) = \text{lv}(b) \cup \text{lv}(b_x) \\
\text{lv}(\text{for } x \text{ in } e : b_x; & b) = x \cup \text{lv}(b) \cup \text{lv}(b_x) \\
\text{lv}(\text{if } e : b_x \text{ else } : b_y; & b) = \text{lv}(b) \cup \text{lv}(b_x) \cup \text{lv}(b_y) \\
\text{lv}(\text{try } : b_x \text{ except } e, x : b_y; & b) = x \cup \text{lv}(b) \cup \text{lv}(b_x) \cup \text{lv}(b_y) \\
\text{lv}(\text{try } : b_x \text{ finally } : b_y; & b) = \text{lv}(b) \cup \text{lv}(b_x) \cup \text{lv}(b_y) \\
\text{lv}(s; & b) = \text{lv}(b)
\end{aligned}$$

Returning from a function

A function stops execution at the end of its block or when a return statement is encountered. When a function returns, it returns a value and the environment stored on the stack is reinstated. Functions always return a value; those without a return statement return `None`.

Returning from the end of a function

When all statements in b_λ have been executed, but no return statement has been executed, we will encounter the return marker $\Gamma \vdash \circ_\lambda$. The return marker contains the old environment Γ

8 Functions and generators

which is reinstated. Since there was no return statement, we return `None`.

$$\begin{aligned} & \langle \Theta, \Gamma_\lambda, S | \Gamma \vdash \circ_\lambda, a \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S, a_{\text{None}} \rangle \end{aligned} \quad (8.7)$$

Executing a return statement

At a return statement, the returned expression e is evaluated and a return frame `return \circ` is placed on the stack.

$$\begin{aligned} & \langle \Theta, \Gamma, S, \text{return } e; \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S | \text{return } \circ, e \rangle \end{aligned} \quad (8.8)$$

Start stack unwinding

Next, once the returned value e has been evaluated to a , we use the return instruction to unwind the stack and return a

$$\begin{aligned} & \langle \Theta, \Gamma, S | \text{return } \circ, a \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S, \text{return } a \rangle \end{aligned} \quad (8.9)$$

Unwinding the stack

Finally the return instruction unwinds (pops frames from) the stack until the return frame is found on top of the stack. At the return frame the old environment is reinstated, a is returned, and unwinding stops.

$$\begin{aligned} & \langle \Theta, \Gamma, S | f, \text{return } a \rangle \\ \Rightarrow & \begin{cases} \langle \Theta, \Gamma', S, a \rangle & \text{if } f = \Gamma' \vdash \circ_\lambda \\ \langle \Theta, \Gamma, S | \text{return } a, b \rangle & \text{if } f = \text{finally} : b \\ \langle \Theta, \Gamma, S, \text{return } a \rangle & \text{otherwise} \end{cases} \end{aligned} \quad (8.10)$$

If the return statement is in the body of a try-finally statement, we execute the finally clause `finally : b` before returning from the function.

8.3 Executing generators

In the previous sections we have already seen how a generator function is defined and executed. In this section we describe how generators are executed.

The execution of a generator function returned a generator $\langle S, \Gamma \rangle$ with a stack S and environment Γ . When the generator itself is executed, the stack S and environment Γ become the stack and environment of the abstract machine. The generator is suspended by storing the stack and environment in the generator value, and reinstating the old stack and environment.

The following rules describe how the generator is (re)started using the primitive functions `__next__` or `__send__`, and how it is suspended by the yield expression.

Getting the next yielded value

The primitive function `__next__` starts or resumes a generator if it hasn't already started.

If `__next__` is called for generator that has already started, i.e., it has a value of $\langle \rangle$, we simply raise a `ValueError`. A running generator cannot be (re)started.

If `__next__` is called for a suspended or new generator $\langle S_{\text{gen}}, \Gamma_{\text{gen}} \rangle$ the stack S_{gen} is pushed on the machine stack and the environment is set to Γ_{gen} . A yield marker $\Gamma \vdash a_{\text{self}}$ with the current environment Γ and the generator's address a_{self} is pushed on the stack, before S_{gen} the generator's stack.

$$\begin{aligned} & \langle \Theta, \Gamma, S, \text{generator}.__\text{next}__([a_{\text{self}}]) \rangle \\ \Rightarrow & \begin{cases} \langle \Theta', \Gamma_{\text{gen}}, S | \Gamma \vdash a_{\text{self}} | S_{\text{gen}}, a_{\text{None}} \rangle & \text{if } v = \langle S_{\text{gen}}, \Gamma_{\text{gen}} \rangle \\ \langle \Theta, \Gamma, S, \text{raise } a_{\text{ValueError}}() \rangle & \text{if } v = \langle \rangle \end{cases} \quad (8.11) \\ \text{where } & \Theta' = \Theta \oplus [a_{\text{self}} \rightarrow \langle \rangle] \\ & v = \Theta(a_{\text{self}}) \end{aligned}$$

Getting the next value with yield value

The primitive `__send__` resumes a generator like `__next__`, but there is one difference: calling `__send__` with argument a causes the yield expression where the generator has suspended, to evaluate to a .

$$\begin{aligned} & \langle \Theta, \Gamma, S, \text{generator}.__\text{send}__([a_{\text{self}}, a]) \rangle \\ \Rightarrow & \begin{cases} \langle \Theta', \Gamma_{\text{gen}}, S | \Gamma \vdash a_{\text{self}} | S_{\text{gen}}, a \rangle & \text{if } v = \langle S_{\text{gen}}, \Gamma_{\text{gen}} \rangle \\ \langle \Theta, \Gamma, S, \text{raise } a_{\text{ValueError}}() \rangle & \text{if } v = \langle \rangle \end{cases} \quad (8.12) \\ \text{where } & \Theta' = \Theta \oplus [a_{\text{self}} \rightarrow \langle \rangle] \\ & v = \Theta(a_{\text{self}}) \end{aligned}$$

The yield expression**Evaluating a yield expression**

First, the value e to be returned is evaluated.

$$\begin{aligned} & \langle \Theta, \Gamma, S, \text{yield } e \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S | \text{yield } \circ, e \rangle \end{aligned} \quad (8.13)$$

Start unwinding of the stack

Then, with the yielded value a , we start to unwind the stack.

$$\begin{aligned} & \langle \Theta, \Gamma, S | \text{yield } \circ, a \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S, \text{yield } a, \square \rangle \end{aligned} \quad (8.14)$$

The second part of the yield instruction $\text{yield } a, \square$ stores the part of the stack that has been unwound. Unwinding starts with an empty stack (\square).

Unwinding the stack

Unwinding of the stack continues until a yield marker $\gamma \vdash a_{\text{gen}}$ is on top. At that point, the generator a_{gen} is updated with the unwound stack S_{gen} and the generator's environment Γ_{gen} . Thus the state of the generator at the yield statement is stored.

$$\begin{aligned} & \langle \Theta, \Gamma, S, \text{yield } a, S_{\text{gen}} \rangle \\ \Rightarrow & \begin{cases} \langle \Theta \oplus [a_{\text{gen}} \rightarrow \langle S_{\text{gen}}, \Gamma_{\text{gen}} \rangle], \Gamma, S, a \rangle & \text{if } f = \Gamma \vdash a_{\text{gen}} \\ \langle \Theta, \Gamma_{\text{gen}}, S, \text{yield } a, f | S_{\text{gen}} \rangle & \text{otherwise} \end{cases} \end{aligned} \quad (8.15)$$

Stopping a generator

A generator stops when it reaches the end of its body, or an exception is raised in its body (see Rule 10.5). Stopped generators cannot be restarted, they will raise a `StopIteration` exception when calling `__next__`.

Stopping at the end of a generator

When the yield marker $\Gamma \vdash a_{\text{gen}}$ is on top of the stack, the generator's body has been executed. We stop the generator, raise a `StopIteration` and reinstate the old environment.

$$\begin{aligned} & \langle \Theta, \Gamma, S | \Gamma \vdash a_{\text{gen}}, a \rangle \\ \Rightarrow & \langle \Theta \oplus [a_{\text{gen}} \rightarrow \langle \square, \square \rangle], \Gamma, S, \text{raise } a_{\text{StopIteration}}() \rangle \end{aligned} \quad (8.16)$$

8.3 Executing generators

The generator stored on the heap at address a_{gen} is updated to the empty generator $\langle \square, \square \rangle$. This will cause subsequent executions of the generator to immediately stop.

8 *Functions and generators*

9 Classes and objects

A number of statements and expressions are devoted to classes and objects. Most of which are common to object-oriented language. Class definitions, attribute access, attribute assignments and class functions are all common programming constructs.

Python differs from most other languages in its extremely dynamic implementation of objects: Python's objects, and even classes, can be altered at any time. Attributes, including the `__class__` and `__bases__` attributes, can be removed and (re)defined using delete and assignment statements. Furthermore, the semantics of attribute references, assignments, and deletions can be changed by overriding class members.

The example in Listing 9.1 shows how the semantics of attribute assignment are overridden by the class function `__setattr__`, and how the attribute deletion statement removes `__setattr__` even after the class has been created.

```
class Counter(object) :
    def __init__(self) :
        object.__setattr__(self, "x", 0)
    def inc(self) :
        object.__setattr__(self, "x", self.x + 1)
    def __setattr__(self, attr, val) :
        print("Use inc() to change the counter value")

c = Counter()
print(c.x)           # prints 0
c.inc()
print(c.x)         # prints 1
c.x = 0             # prints the error message
print(c.x)         # prints 1

del Counter.__setattr__ # remove the class function __setattr__
c.x = 5              # nothing is printed
print(c.x)         # prints 5
```

Listing 9.1: A simple use-case for classes and objects. The Counter class overrides `__setattr__` to prevent direct access to the internal counter `x`. Because the class overrides `__setattr__` we are forced to use `object.__setattr__` to set the attributes of a counter.

9.1 Class definition

A class definition `class x(e) : b` creates a new class object. Its bases are set to the value of `e`, and the variables bound in `b` define its attributes.

9 Classes and objects

Class definitions are not declarative, like in many other object oriented languages, but imperative in style. The body b and bases e of a new class are executed and evaluated when the class definition is executed.

Evaluating the bases of a new class

Before the body is executed, the base objects have to be evaluated. The bases are represented by a tuple expression e .

$$\Rightarrow \langle \Theta, \Gamma, S, \text{class } x(e) : b; \rangle$$

$$\Rightarrow \langle \Theta, \Gamma, S | \text{class } x(o) : b, e \rangle \quad (9.1)$$

Although one would expect to derive only from instances of `type`, any object is accepted as a base at this point.

Executing the body of a class definition

When e has been evaluated, its result a_b is stored on the stack.

The body is executed with a new object γ_c on top of the environment stack. All variable bindings within the body will be bound in the object γ_c . Thus attributes of the new class are defined by ordinary variable bindings in the body.

$$\Rightarrow \langle \Theta, \Gamma | \gamma_1, S | \text{class } x(o) : b, a_b \rangle$$

$$\Rightarrow \langle \Theta \oplus [\gamma_c \rightarrow (\emptyset, \text{None})], \Gamma', S | \Gamma | \gamma_1 \vdash \text{class } x(a_b) : o, b \rangle$$

where $\gamma_c = \text{new address} \notin \Theta$ (9.2)

$$\Gamma' = \text{if } \Theta(\gamma_1) \text{ is object then } \Gamma | \gamma_c$$

$$\text{else } \Gamma | \gamma_1 | \gamma_c$$

Class definitions that are nested in class definitions, can not access the attributes of the class enclosing them through the environment. If the topmost environment γ_1 is an object, the body of the class definition is executed without γ_1 in the environment stack. Nested function definitions are limited in the same way (see Equation 8.1).

Creating a new class

After the body has been executed, we store the class name x on the heap and construct a new class object. The new object is constructed by the constructor of the metaclass i.e., the class function `--new--`. See Rule 9.4 for the default constructor.

$$\Rightarrow \langle \Theta, \Gamma | \gamma_c, S | \Gamma | \gamma_1 \vdash \text{class } x(a_b) : o, a \rangle$$

$$\Rightarrow \langle \Theta', \Gamma | \gamma_1, S | o_x \text{--init--}(a_x, a_b, \gamma_c), \Phi_{a_t}^\Theta(\text{--new--})([a_t, a_x, a_b, \gamma_c]) \rangle$$

where $\Theta' = \Theta \oplus [a_x \rightarrow x]$ (9.3)

$$a_x = \text{new address} \notin \Theta$$

$$a_t = \text{if } \text{--metaclass--} \in \Theta(\gamma_c) \text{ then } \Theta(\gamma_c)(\text{--metaclass--})$$

$$\text{else } a_{\text{type}}$$

Note that this implementation of metaclasses is incomplete: in Python the metaclass is inherited from the base(s) of a class. This is more complicated than it seems at first glance. For example, when a new class C is created that inherits from two classes A and B which inherit from two different metaclasses M_A and M_B respectively, it is unclear which metaclass constructs C . Python will raise an exception in this particular situation. For the moment *minpy* avoids these complications by disabling inheritance of metaclasses.

The default class constructor

The default class constructor extends the new class a_c with name a_x , bases a_b and class a_{self} . The `__metaclass__` attribute, which typically refers to a_{self} , is removed from the new class. If no valid inheritance graph can be created because the bases are bad (see 5.1), we raise a `TypeError`.

$$\begin{aligned}
 & \langle \Theta, \Gamma, S, \text{type}.__\text{new}__([a_{\text{self}}, a_x, a_b, a_c]) \rangle \\
 \Rightarrow & \begin{cases} \langle \Theta', \Gamma, S, a_c & \rangle \text{ if } \Theta'(a_c) \text{ is valid} \\ \langle \Theta, \Gamma, S, \text{raise } a_{\text{TypeError}}() & \rangle \text{ otherwise} \end{cases} \\
 & \text{where } \Theta' = \Theta \oplus [a_c \rightarrow (\Theta(a_c) \oplus [__\text{class}__ \rightarrow a_{\text{self}} \\
 & \hspace{10em}, __\text{bases}__ \rightarrow a_b \\
 & \hspace{10em}]) \\
 & \hspace{10em} \ominus __\text{metaclass}__ \\
 & \hspace{10em}]
 \end{aligned} \tag{9.4}$$

Initializing the new class

New classes, created by the constructor `__new__` of the metaclass, are initialized by the `__init__` function of the metaclass. So when the constructor returns, the initializer is called.

The assignment on the stack binds the class name x to the new class a_o . If the initialization function raises an exception, the assignment will be popped from the stack and the class will not be bound.

$$\begin{aligned}
 & \langle \Theta, \Gamma, S | \circ_x __\text{init}__(a_x, a_b, \gamma_c), a_o \rangle \\
 \Rightarrow & \langle \Theta, \Gamma, S | x = a_o; \hspace{10em}, \Phi_{a_t}^\ominus(__\text{init}__)([a_t, a_x, a_b, \gamma_c]) \rangle \\
 & \text{where } a_t = \Theta(a_o)(__\text{class}__)
 \end{aligned} \tag{9.5}$$

The original metaclass, whose constructor was called, can return any object. Therefore we call the initialization function of the returned object's class, rather than the original metaclass.

9.2 Object creation

The creation of objects is handled by the `__call__`, `__new__`, and `__init__` attributes. Most classes will inherit those attributes from the base object `object`, whose primitives define the default construction mechanism.

The default constructor

Every class can be called as if it were a function because all classes inherit, or override, this `__call__` primitive. If a class is called, its `__call__` attribute function is called (see Rule 8.6). The `__call__` primitive consecutively calls the `__new__` and `__init__` to construct a new class and initialize it.

$$\begin{aligned} & \langle \Theta, \Gamma, S, \text{type.}__call__(a_c : \bar{a}) \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S \mid \circ _init(\bar{a}), a_f(a_c : \bar{a}) \rangle \\ \text{where } a_f = & (\bigoplus_{\Theta} \text{MRO } \Theta(a_c))(_new) \end{aligned} \quad (9.6)$$

The `__new__` function a_f , is looked up in the class itself or its bases. The frame $\circ _init(\bar{a})$, is put onto the stack to execute the `__init__` with the same arguments.

The default object creation

The default implementation of `__new__` in `object` creates a new object with base `None`, and sets its class to a_c .

$$\begin{aligned} & \langle \Theta, \Gamma, S, \text{object.}__new__(a_c : \bar{a}) \rangle \\ \Rightarrow & \langle \Theta \oplus [a \rightarrow ([_class \rightarrow a_c], \text{None})], \Gamma, S, a \rangle \\ \text{where } a = & \text{new address} \notin \Theta \end{aligned} \quad (9.7)$$

Primitive values that allow subclassing, have their own `__new__` attributes which initialize the base value to their primitive value. See for example Rule A.8.

From creation to initialization

The new object a returned by `__new__` is initialized by its class. If the initialization function does not raise an exception, the new object a on the stack will be returned.

$$\begin{aligned} & \langle \Theta, \Gamma, S \mid \circ _init(\bar{a}), a \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S \mid a, a_f(a : \bar{a}) \rangle \\ \text{where } a_f = & \Phi_a^\Theta(_init) \end{aligned} \quad (9.8)$$

The default object initialization

The default initialization does nothing at all and returns `None`.

$$\begin{aligned} & \langle \Theta, \Gamma, S, \text{object.}__init__(\bar{a}) \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S, a_{\text{None}} \rangle \end{aligned} \quad (9.9)$$

Like `object.__new__`, this primitive accepts an arbitrary number of arguments. This allows a user defined `__new__` function to accept any number of arguments.

9.3 Attribute access

The attribute access expression $e.x$ evaluates to the attribute x of object e . However, its semantics can be changed by overriding the `__getattribute__` function.

Evaluating the attribute access expression

First the left hand side expression e of the expression $e.x$ is evaluated.

$$\begin{aligned} & \langle \Theta, \Gamma, S, e.x \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S | \circ .x, e \rangle \end{aligned} \quad (9.10)$$

Calling `__getattribute__`

When the left hand side expression e has been evaluated to an object a , the `__getattribute__` function of a is called. The function `__getattribute__` is called with the evaluated object a (we instantiate `__getattribute__` with a) and the attribute name a_x .

$$\begin{aligned} & \langle \Theta, \Gamma, S | \circ .x, a \rangle \\ \Rightarrow & \langle \Theta \oplus [a_x \rightarrow x], \Gamma, S, a_f([a, a_x]) \rangle \\ \text{where } & a_f = \Phi_a^\Theta(\text{__getattribute__}) \\ & a_x = \text{new address} \notin \Theta \end{aligned} \quad (9.11)$$

The `__getattribute__` function is retrieved from the class of the object a . This means that it cannot be overridden by instance attributes.

Getting attributes and instantiating function attributes

The default implementation of `__getattribute__` for objects retrieves an attribute according to the MRO (see Section 5.1). When a class function of an object o is retrieved it is automatically instantiated with the object, i.e., its first argument is set to o . Once instantiated, a class function will not be instantiated again.

We have four cases:

1. If a_x is an instance attribute, we return its value uninstantiated.
2. If a_x is an inherited attribute and is instantiatable, we return the instantiated attribute.
3. If a_x is an inherited attribute but is not instantiatable, we return it.
4. Otherwise the attribute is not defined so we raise an `AttributeError`.

$$\begin{aligned}
 & \langle \Theta, \Gamma, S, \text{object}._\text{getattribute}__([a_o, a_x]) \rangle \\
 \Rightarrow & \left\{ \begin{array}{l} \langle \Theta, \Gamma, S, a \rangle \\ \langle \Theta, \Gamma, S, a_o.a \rangle \\ \langle \Theta, \Gamma, S, a \rangle \\ \langle \Theta, \Gamma, S, \text{raise } a_{\text{AttributeError}}() \rangle \end{array} \right. \begin{array}{l} \text{if } x \in \Omega_{a_o}^\Theta \\ \text{if } x \notin \Omega_{a_o}^\Theta \wedge x \in \Phi_{a_o}^\Theta \wedge \Theta(a) \text{ is instantiable} \\ \text{if } x \notin \Omega_{a_o}^\Theta \wedge x \in \Phi_{a_o}^\Theta \wedge \Theta(a) \text{ is not instantiable} \\ \text{otherwise} \end{array} \\
 & \text{where } a = \Upsilon_{a_o}^\Theta(x) \\
 & \quad x = \Theta(a_x)
 \end{aligned} \tag{9.12}$$

Note that only uninstantiated functions are instantiable. Callable objects are instantiated when it is called (see Rule 8.6).

Executing an instantiated function

A function a_f instantiated with object a_o is denoted as $a_o.a_f$. When an instantiated function is evaluated in a function application, the object a_o is inserted as the first argument of the function call.

$$\begin{aligned}
 & \langle \Theta, \Gamma, S, a_o.a_f(\bar{a}) \rangle \\
 \Rightarrow & \langle \Theta, \Gamma, S, a_f(a_o : \bar{a}) \rangle
 \end{aligned} \tag{9.13}$$

Getting class members and inserting type checks

The `__getattribute__` primitive of `type` is equivalent to the `__getattribute__` primitive of `object`, except for one difference: we insert type annotations for instantiatable class functions. The type annotations ensure that class functions are only applied to instances of their class.

$$\begin{aligned}
 & \langle \Theta, \Gamma, S, \text{type}._\text{getattribute}__([a_o, a_x]) \rangle \\
 \Rightarrow & \left\{ \begin{array}{l} \langle \Theta, \Gamma, S, a : [a_o] \rangle \\ \langle \Theta, \Gamma, S, \text{object}._\text{getattribute}__([a_o, a_x]) \rangle \end{array} \right. \begin{array}{l} \text{if } x \in \Omega_{a_o}^\Theta \wedge \Theta(a) \text{ is instantiable} \\ \text{otherwise} \end{array} \\
 & \text{where } a = \Upsilon_{a_o}^\Theta(x) \\
 & \quad x = \Theta(a_x)
 \end{aligned} \tag{9.14}$$

9.4 Attribute assignment

By default, the attribute assignment statement $e_o.x = e$ (re)defines the attribute x of the target object e_o to the value e . The semantics of attribute assignment statements are determined by the `__setattr__` attribute of an object's class. By overloading it, a programmer can redefine the semantics.

Evaluating the right hand side of an attribute assignment

The evaluation of the assignment starts with the right hand side: the value e . Note that this behaviour is different from the usual left to right evaluation.

$$\begin{aligned} & \langle \Theta, \Gamma, S, e_o.x = e; \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S | e_o.x = \circ, e \rangle \end{aligned} \quad (9.15)$$

Evaluating the left hand side of an attribute assignment

Next, the left hand side e_o is evaluated, while the value a of the right hand side is placed on the stack.

$$\begin{aligned} & \langle \Theta, \Gamma, S | e_o.x = \circ, a \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S | \circ.x = a, e_o \rangle \end{aligned} \quad (9.16)$$

Calling `__setattr__`

Finally we call the inherited attribute `__setattr__` of a_o , with the target object a_o , the assigned attribute name a_x , and the new value a .

$$\begin{aligned} & \langle \Theta, \Gamma, S | \circ.x = a, a_o \rangle \\ \Rightarrow & \langle \Theta \oplus [a_x \rightarrow x], \Gamma, S, a_f([a_o, a_x, a]) \rangle \\ \text{where } & a_f = \Phi_{a_o}^{\Theta}(\text{__setattr__}) \\ & a_x = \text{new address} \notin \Theta \end{aligned} \quad (9.17)$$

The default implementation of `__setattr__` for objects

The default implementation of `__setattr__` for objects updates the object a_o with $[x \rightarrow a]$.

$$\begin{aligned} & \langle \Theta, \Gamma, S, \text{object.__setattr__}([a_o, a_x, a]) \rangle \\ \Rightarrow & \langle \Theta \oplus [a_o \rightarrow o'], \Gamma, S, a_{\text{None}} \rangle \\ \text{where } & o' = \Theta(a_o) \oplus [\Theta(a_x) \rightarrow a] \end{aligned} \quad (9.18)$$

9.5 Attribute deletion

By default, an attribute deletion statement `del e_o.x` removes the attribute x from object e_o if it exists. The semantics can be redefined by overloading `__delattr__`.

Evaluating the left hand side of an attribute deletion

We start by evaluating the object e_o .

$$\begin{aligned} & \langle \Theta, \Gamma, S, \text{del } e_o.x; \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S | \text{del } \circ.x, e_o \rangle \end{aligned} \quad (9.19)$$

Calling `__delattr__`

When e_o has been evaluated to a_o , we call the `__delattr__` attribute inherited by a_o .

$$\begin{aligned} & \langle \Theta, \Gamma, S | \text{del } \circ.x, a_o \rangle \\ \Rightarrow & \langle \Theta \oplus [a_x \rightarrow x], \Gamma, S, a_f([a_o, a_x]) \rangle \\ \text{where } & a_f = \Phi_{a_o}^{\Theta}(\text{__delattr__}) \\ & a_x = \text{new address} \notin \Theta \end{aligned} \quad (9.20)$$

The default implementation of `__delattr__`

The default implementation of `__delattr__` removes the attribute a_x from the object a_o if x is defined in a_o . Otherwise an `AttributeError` is raised.

$$\begin{aligned} & \langle \Theta, \Gamma, S, \text{object.__delattr__}([a_o, a_x]) \rangle \\ \Rightarrow & \begin{cases} \langle \Theta \oplus [a_o \rightarrow o'], \Gamma, S, a_{\text{None}} \rangle & \text{if } \Theta(a_x) \in \Theta(a_o) \\ \langle \Theta, \Gamma, S, \text{raise } a_{\text{AttributeError}}() \rangle & \text{otherwise} \end{cases} \\ \text{where } & o' = \Theta(a_o) \ominus \Theta(a_x) \end{aligned} \quad (9.21)$$

10 Exceptions

Exceptions simplify handling of errors by providing an alternative mode of execution. When an exceptional situation occurs i.e., an exception is raised, normal execution stops, until the situation is dealt with.

Listing 10.1 shows how exceptions can be used to handle divisions by zero, how normal execution stops when an exception is raised, and how the exception is handled.

```
class ZeroDivisionError(Exception) :
    def __str__(self) :
        return "Divided a number by zero!"

def div(x, y) :
    if y == 0 :
        raise ZeroDivisionError()

    return x / y

try :
    div(2, 0) # will raise the exception
    print("This will never be printed")
except Exception, e : # catch all exceptions from within the try block
    print(e)

print("Normal execution continues")
```

Listing 10.1: Raising and catching a divide by zero exception

10.1 The raise statement

A raise statement `raise e` stops normal execution and raises `e` if and only if it evaluates to an instance of `BaseException`. During normal execution we would continue with the next frame on the stack, whereas most stack frames are simply discarded while raising an exception. Only some stack frames and markers will be handled while raising an exception.

Executing the raise statement

The raise statement first evaluates the exception expression `e`.

$$\Rightarrow \langle \Theta, \Gamma, S, \text{raise } e; \rangle \quad (10.1)$$

Start unwinding of the stack

If the raised value a is an instance of `BaseException`, unwinding of the stack starts. Otherwise, a `TypeError` is raised. Raising an exception has no effect on the environment.

$$\begin{aligned} & \langle \Theta, \Gamma, S | \text{raise } \circ, a \rangle \\ \Rightarrow & \begin{cases} \langle \Theta, \Gamma, S & , \text{raise } a \rangle & \text{if } (a :_{\Theta} a_{\text{BaseException}}) \\ \langle \Theta, \Gamma, S & , \text{raise } a_{\text{TypeError}}() \rangle & \text{otherwise} \end{cases} \end{aligned} \quad (10.2)$$

Raising an exception

By default, stack frames are simply discarded while unwinding the stack until the stack is empty and execution stops. Note that this is a very unspecific rule and hence will be overridden by rules such as Rule 10.4.

$$\begin{aligned} & \langle \Theta, \Gamma, S | f, \text{raise } a \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S & , \text{raise } a \rangle \end{aligned} \quad (10.3)$$

Raising an exception in a function

If an exception remains uncaught in a function, we return from the function by reinstating the environment Γ' found on the return marker introduced in Rule 8.6, and resume unwinding.

$$\begin{aligned} & \langle \Theta, \Gamma, S | \Gamma' \vdash \circ_{\lambda}, \text{raise } a \rangle \\ \Rightarrow & \langle \Theta, \Gamma', S & , \text{raise } a \rangle \end{aligned} \quad (10.4)$$

Raising an exception in a generator

When an exception is raised in a generator, we will encounter the yield marker $\Gamma' \vdash a_g$ introduced in Rule 8.11. Once a generator has raised an exception, it is stopped: the current generator a_g is updated with an empty generator $\langle \square, \square \rangle$ and the old environment Γ' is reinstated (see also Rule 8.16).

$$\begin{aligned} & \langle \Theta & , \Gamma, S | \Gamma' \vdash a_g, \text{raise } a \rangle \\ \Rightarrow & \langle \Theta \oplus [a_g \rightarrow \langle \square, \square \rangle], \Gamma', S & , \text{raise } a \rangle \end{aligned} \quad (10.5)$$

10.2 The try-except statement

A raised exception is handled by a try-except statement. A try-except statement `try : b except e, x : be` executes its body *b*. If an exception is raised in the body and it is an instance of *e*, the raised exception is bound to *x* and normal execution resumes with the except clause *b_e*. If a raised exception is not an instance of *e*, the exception is re-raised and unwinding continues.

Executing the try-except statement

The try-except statement places the except clause `except e, x : be` on the stack, and executes the body *b*.

$$\begin{aligned} & \langle \Theta, \Gamma, S, \text{try} : b \text{ except } e, x : b_e; \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S | \text{except } e, x : b_e, b \rangle \end{aligned} \quad (10.6)$$

Popping an exception clause while unwinding

When an except clause is encountered while unwinding, its exception class *e* is evaluated in order to check whether it catches the exception. The raised exception *a* is put on the stack, along with the exception clause.

$$\begin{aligned} & \langle \Theta, \Gamma, S | \text{except } e, x : b_e, \text{raise } a \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S | \overset{\text{raise } a}{\text{except } e, x : b_e}, e \rangle \end{aligned} \quad (10.7)$$

Catching an exception

Once the class of the except clause has been evaluated, we determine whether it catches the raised exception:

1. If the class *a_e* of the except clause is an exception (i.e. it is a `BaseException`) and the raise exception *a* is an instance of *a_e*, we catch the exception. The exception *a* is bound to *x* and the *b_e* is executed.
2. Otherwise, we re-raise the exception *a* and pop the except clause.

$$\begin{aligned} & \langle \Theta, \Gamma | \gamma_1, S | \overset{\text{raise } a}{\text{except } e, x : b_e}, a_e \rangle \\ \Rightarrow & \begin{cases} \langle \Theta', \Gamma | \gamma_1, S, b_e \rangle & \text{if } a_e \leq_{\Theta} a_{\text{BaseException}} \wedge a :_{\Theta} a_e \\ \langle \Theta, \Gamma | \gamma_1, S, \text{raise } a \rangle & \text{otherwise} \end{cases} \quad (10.8) \\ & \text{where } \Theta' = \Theta \oplus [\gamma_1 \rightarrow \Theta(\gamma_1) \oplus [x \rightarrow a]] \end{aligned}$$

The class of an except clause *a_e* can be any value. It is not at all limited to classes, or even exceptions. However, if *a_e* is not an exception class, it will never catch anything.

Popping the exception marker

The exception clause is simply popped from the stack when no exception has been raised in the body.

$$\begin{aligned} & \langle \Theta, \Gamma, S | \text{except } e, x : b_e, a \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S \quad \quad \quad , a \rangle \end{aligned} \tag{10.9}$$

10.3 The try-finally statement

Sometimes it is necessary to perform some operations even when an error has occurred. The try-finally statement `try : b finally : bf` executes its body *b* with a guarantee that *b_f* will be executed. We would typically use a try-finally statement to ensure some clean-up code is executed, even if an exception is raised in its body.

The guarantee of executing *b_f* applies not only to exception handling. For example, if there is a return statement in the body *b*, we will execute *b_f* before we return from the function. If *b_f* contains another return statement, the last return statement will return from the function. This rather unintuitive behaviour has prompted the CPython developers to give a warning when the *b_f* contains a yield or return statement.

Executing the try-finally statement

The try-finally starts like the try-except statement, by executing its body. The finally block *b_f* is placed on the stack.

$$\begin{aligned} & \langle \Theta, \Gamma, S \quad \quad \quad , \text{try} : b \text{ finally} : b_f; \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S | \text{finally} : b_f, b \quad \quad \quad \rangle \end{aligned} \tag{10.10}$$

Popping a finally clause while unwinding

When a finally clause is encountered while unwinding, the raised exception *a* is put on the stack and the finally clause *b_f* is executed.

$$\begin{aligned} & \langle \Theta, \Gamma, S | \text{finally} : b_f, \text{raise } a_e \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S | \text{raise } a_e \quad \quad \quad , b_f \quad \quad \quad \rangle \end{aligned} \tag{10.11}$$

Popping the finally marker

When no exception has occurred, the try-finally statement executes the finally clause that has been placed on the stack.

$$\Rightarrow \langle \Theta, \Gamma, S | \text{finally} : b_f, a \rangle \Rightarrow \langle \Theta, \Gamma, S, b_f \rangle \quad (10.12)$$

10 *Exceptions*

11 Control structures

Python includes three basic control structures: `if`, `while` and `for` statements. All three are very common in imperative programming languages and have similar semantics in Python. The `if` statement executes one of its two blocks of statements, depending on the value of a test expression. The `while` expression repeats its body as long as its test expression is nonzero. The `for` statement iterates over a number of values, executing its body once for every value.

Unlike many other languages, the control structures of Python are non-scoping compound statements, i.e., statements that contain blocks of statements but do not define a lexical scope (see Chapter 3.2). This means that the variables assigned inside an `if`, `while` or `for` statement are still bound after the execution of the statements.

The `break` and `continue` statements, respectively stop or interrupt a loop. They can only occur in the body of a `for` or `while` loop.

11.1 The `if` statement

An `if` statement `if e : bif else : belse` consists of the test expression e , the blocks b_{if} and b_{else} . If e evaluates to a nonzero value, b_{if} is executed, else b_{else} is executed. We consider a value to be nonzero, if its class function `__nonzero__` returns `True`.

Evaluating the test expression of an `if` statement

First, the test expression e is evaluated, while the frame `if ○ .__nonzero__() : bif else : belse` indicates that we still need to call `__nonzero__` and execute either b_{if} or b_{else} .

$$\Rightarrow \left\langle \Theta, \Gamma, S \right\rangle, \text{if } e : b_{\text{if}} \text{ else : } b_{\text{else}} ; \left. \right\rangle \quad (11.1)$$
$$\Rightarrow \left\langle \Theta, \Gamma, S \mid \text{if } \circ .\text{__nonzero__() : } b_{\text{if}} \text{ else : } b_{\text{else}}, e \right\rangle$$

Testing the test value of an `if` statement

Next, we call the class function `__nonzero__` to determine which block should be executed. The stack frame `if ○ : bif else : belse` indicates that we are evaluating $a.\text{__nonzero__}()$ and contains the two blocks. If the test value a has no class attribute `__nonzero__`, we raise a `TypeError`.

$$\begin{aligned}
& \langle \Theta, \Gamma, S | \text{if } \circ \text{_nonzero_}() : b_{\text{if}} \text{ else } : b_{\text{else}}, a \rangle \\
\Rightarrow & \begin{cases} \langle \Theta, \Gamma, S | \text{if } \circ : b_{\text{if}} \text{ else } : b_{\text{else}} & , a_f([a]) \rangle & \text{if } \text{_nonzero_} \in \Phi_a^\Theta \\ \langle \Theta, \Gamma, S & , \text{raise } a_{\text{TE}}() \rangle & \text{otherwise} \end{cases} \\
& \text{where } a_f = \Phi_a^\Theta(\text{_nonzero_}) \\
& \quad a_{\text{TE}} = a_{\text{TypeError}}
\end{aligned} \tag{11.2}$$

Executing the body of an if statement

Depending on the value of the test expression we execute either the **if** or the **else** block:

1. If `__nonzero__` returned a boolean, and it is `True`, we execute `bif`.
2. If `__nonzero__` returned a boolean, and it is `False`, we execute `belse`.
3. Otherwise, the call to `__nonzero__` returned a different class, so we return a `TypeError`.

$$\begin{aligned}
& \langle \Theta, \Gamma, S | \text{if } \circ : b_{\text{if}} \text{ else } : b_{\text{else}}, a \rangle \\
\Rightarrow & \begin{cases} \langle \Theta, \Gamma, S & , b_{\text{if}} \rangle & \text{if } a : \Theta_{a_{\text{bool}}} \wedge \Theta(a) \\ \langle \Theta, \Gamma, S & , b_{\text{else}} \rangle & \text{if } a : \Theta_{a_{\text{bool}}} \wedge \neg \Theta(a) \\ \langle \Theta, \Gamma, S & , \text{raise } a_{\text{TypeError}}() \rangle & \text{otherwise} \end{cases} \tag{11.3}
\end{aligned}$$

11.2 The while statement

The while statement `while e : b` consists of a test expression `e` and a body `b`. Like the `if` statement, the while statement depends on the value of the test expression: if it evaluates to a nonzero value, the body is executed and the while statement (including the test) is executed again, otherwise the while-statement ends.

Evaluating the test expression of a while statement

First, we evaluate the test expression `e` and place the while statement on the stack. The frame `while \circ_e __nonzero__() : b` indicates that we still have to call `__nonzero__` and contains both the body and test expression.

$$\begin{aligned}
& \langle \Theta, \Gamma, S & , \text{while } e : b; \rangle \\
\Rightarrow & \langle \Theta, \Gamma, S | \text{while } \circ_e \text{_nonzero_}() : b, e \rangle
\end{aligned} \tag{11.4}$$

Testing the test value of a while statement

Next, we call the class function `__nonzero__` of the test value a , or, if it has no such attribute, raise a `TypeError`.

$$\begin{aligned} & \langle \Theta, \Gamma, S | \text{while } \circ_e \text{__nonzero__}() : b, a \rangle \\ \Rightarrow & \begin{cases} \langle \Theta, \Gamma, S | \text{while } \circ_e : b & , a_f([a]) \rangle & \text{if } \text{__nonzero__} \in \Phi_a^\Theta \\ \langle \Theta, \Gamma, S & , \text{raise } a_{\text{TypeError}}() \rangle & \text{otherwise} \end{cases} \quad (11.5) \\ & \text{where } a_f = \Phi_a^\Theta(\text{__nonzero__}) \end{aligned}$$

Executing the body of a while statement

We execute the body b and rerun the while loop, if the test expression evaluated to a nonzero value. If not, we stop the loop:

1. If `__nonzero__` returned a boolean, and it is `True`, we execute the body, and place a while marker `while e : \circ_b` on the stack.
2. If `__nonzero__` returned a boolean, and it is `False`, we stop of the while statement and return `None`.
3. Otherwise, the call to `__nonzero__` returned a different class, so we return a `TypeError`.

$$\begin{aligned} & \langle \Theta, \Gamma, S | \text{while } \circ_e : b, a \rangle \\ \Rightarrow & \begin{cases} \langle \Theta, \Gamma, S | \text{while } e : \circ_b, b \rangle & \text{if } a : \Theta a_{\text{bool}} \wedge \Theta(a) \\ \langle \Theta, \Gamma, S & , a_{\text{None}} \rangle & \text{if } a : \Theta a_{\text{bool}} \wedge \neg \Theta(a) \\ \langle \Theta, \Gamma, S & , \text{raise } a_{\text{TypeError}}() \rangle & \text{otherwise} \end{cases} \quad (11.6) \end{aligned}$$

Repeating a while statement

After the execution of the while body, the while marker on top of the stack is executed again.

$$\Rightarrow \langle \Theta, \Gamma, S | \text{while } e : \circ_b, a \rangle \Rightarrow \langle \Theta, \Gamma, S & , \text{while } e : b; \rangle \quad (11.7)$$

One might think that the while marker can be replaced by a simple while statement on the stack. However, the rules that handle `continue` and `break` statements depend on the difference (see Sections 11.4 and 11.5).

11.3 The for statement

A for statement `for x in e : b` consists of a variable name x to which the values produced by the iterable e are bound, and the body b .

An iterable is any object with a class function `__iter__` that produces an iterator: an object with a class function `__next__` that returns the next value if there is one, and raises a `StopIteration` exception when not. In particular, every generator is an iterator and every generator function is an iterable.

A for statement first creates an iterator by calling the class function `__iter__` of the iterable e . Then it executes the body b repeatedly, calling the `__next__` function of the iterator. Binding the variable x to the value returned by `__next__`. Once the `__next__` function raises a `StopIteration` exception, the for statement stops.

Executing the for statement

First, the iterable e is evaluated and the rest of the for statement is put on the stack. The stack frame `for x in e .__iter__() : b` indicates that we need to call e .`__iter__`.

$$\begin{aligned} & \langle \Theta, \Gamma, S, \text{for } x \text{ in } e : b; \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S | \text{for } x \text{ in } e . _ _ \text{iter} _ _ () : b, e \rangle \end{aligned} \quad (11.8)$$

Get the iterator of the iterable

Next, if the iterable has an `__iter__` attribute, we call it, otherwise we raise a `TypeError`.

$$\begin{aligned} & \langle \Theta, \Gamma, S | \text{for } x \text{ in } e . _ _ \text{iter} _ _ () : b, a \rangle \\ \Rightarrow & \left\{ \begin{array}{l} \langle \Theta, \Gamma, S | \text{for } x \text{ in } e : b, a_f([a]) \rangle \text{ if } _ _ \text{iter} _ _ \in \Phi_a^\Theta \\ \langle \Theta, \Gamma, S, \text{raise } a_{\text{TypeError}}() \rangle \text{ otherwise} \end{array} \right. \quad (11.9) \\ & \text{where } a_f = \Phi_a^\Theta(_ _ \text{iter} _ _) \end{aligned}$$

Start the for loop with the iterator

Then, we start the for loop with the iterator a_i , that has been returned by `__iter__`, and push a for marker `for x in a_i : b` on the stack to start an iteration.

$$\begin{aligned} & \langle \Theta, \Gamma, S | \text{for } x \text{ in } e : b, a_i \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S | \text{for } x \text{ in } a_i : b, a_{\text{None}} \rangle \end{aligned} \quad (11.10)$$

The notation of the for marker suggests that the body is being executed. This is clearly not the case, instead, we pretend to just have executed the body in order to trigger Rule 11.11.

Starting an iteration of the for loop

An iteration starts with a call to the `__next__` function attribute of the iterator a_i . If no such attribute is available, then a_i is no iterator, so a `TypeError` is raised.

$$\begin{aligned} & \langle \Theta, \Gamma, S | \text{for } x \text{ in } a_i : \circ_b, a \rangle \\ \Rightarrow & \begin{cases} \langle \Theta, \Gamma, S | \text{for } x \text{ in } a_i : b, a_f([a_i]) \rangle & \text{if } __next__ \in \Phi_{a_i}^\Theta \\ \langle \Theta, \Gamma, S, \text{raise } a_{\text{TypeError}}() \rangle & \text{otherwise} \end{cases} \\ & \text{where } a_f = \Phi_{a_i}^\Theta(__next__) \end{aligned} \quad (11.11)$$

Executing the body of a for loop

The value a returned by the iterator, is bound to x and the body b is executed. We schedule the next iteration by pushing a for marker `for x in $a_i : \circ_b$` on the stack.

$$\begin{aligned} & \langle \Theta, \Gamma | \gamma_1, S | \text{for } x \text{ in } a_i : b, a \rangle \\ \Rightarrow & \langle \Theta \oplus [\gamma_1 \rightarrow m], \Gamma | \gamma_1, S | \text{for } x \text{ in } a_i : \circ_b, b \rangle \\ & \text{where } m = \Theta(\gamma_1) \oplus [x \rightarrow a] \end{aligned} \quad (11.12)$$

Stopping the for loop

The for loop stops, when the iterator's `__next__` function attribute raises a `StopIteration`: If we encounter a stack frame `for x in $a_i : b$` while unwinding the stack with an exception a , we check whether the exception a is a `StopIteration`. If it is, normal execution resumes and we return `None`. Otherwise, the frame is popped and unwinding continues.

$$\begin{aligned} & \langle \Theta, \Gamma, S | \text{for } x \text{ in } a_i : b, \text{raise } a \rangle \\ \Rightarrow & \begin{cases} \langle \Theta, \Gamma, S, a_{\text{None}} \rangle & \text{if } a : \Theta a_{\text{StopIteration}} \\ \langle \Theta, \Gamma, S, \text{raise } a \rangle & \text{otherwise} \end{cases} \end{aligned} \quad (11.13)$$

11.4 The break statement

A break statement immediately terminates execution of the current block and continues execution of the block that encapsulates the while or for statement. In other words, a break statement jumps directly to the end of the loop.

Rewinding the stack for a break statement

Break statements unwind the stack until a frame indicating the end of a while or for loop is encountered. If we encounter a finally clause `finally : b` (see Section 10.10), we execute its

11 Control structures

body b and put the break statement on the stack. When b has been executed, unwinding will continue with the execution of the break statement.

$$\Rightarrow \left\langle \begin{array}{l} \langle \Theta, \Gamma, S | f \quad , \text{break}; \rangle \\ \langle \Theta, \Gamma, S \quad , a_{\text{None}} \rangle \text{ if } f = \text{while } e : \circ_b \\ \langle \Theta, \Gamma, S \quad , a_{\text{None}} \rangle \text{ if } f = \text{for } x \text{ in } a_i : \circ_b \\ \langle \Theta, \Gamma, S | \text{break}; , b \rangle \text{ if } f = \text{finally} : b \\ \langle \Theta, \Gamma, S \quad , \text{break}; \rangle \text{ otherwise} \end{array} \right\rangle \quad (11.14)$$

11.5 The continue statement

Like a break statement, a continue statement terminates execution of a loop's body. But rather than jumping to the end of the loop, it starts another iteration.

Rewinding the stack for a continue statement

Continue statements unwind the stack until the end of a for or while loop. At the end of a loop, the loop is executed again. Like with break statements, any finally clauses will be executed 'on the way out'.

In case of a for or while loop, the for or while marker is left on the stack, and normal execution is resumed. This will cause Rule 11.7 or Rule 11.11 respectively, to run another iteration.

$$\Rightarrow \left\langle \begin{array}{l} \langle \Theta, \Gamma, S | f \quad , \text{continue}; \rangle \\ \langle \Theta, \Gamma, S | \text{while } e : \circ_b \quad , a_{\text{None}} \rangle \text{ if } f = \text{while } e : \circ_b \\ \langle \Theta, \Gamma, S | \text{for } x \text{ in } a_i : \circ_b , a_{\text{None}} \rangle \text{ if } f = \text{for } x \text{ in } a_i : \circ_b \\ \langle \Theta, \Gamma, S | \text{continue}; \quad , b \rangle \text{ if } f = \text{finally} : b \\ \langle \Theta, \Gamma, S \quad , \text{continue}; \rangle \text{ otherwise} \end{array} \right\rangle \quad (11.15)$$

12 Print statements

The print statement `print(e)` prints a string representation of *e* to the standard output. The string representation of *e* is determined by its class function `__str__`. The program in Listing 12.1 defines an object with a special string representation.

```

class Person(object) :
    def __init__(self, name, age) :
        self.name = name
        self.age = age
    def __str__(self) :
        return self.name + " is " + self.age.__str__() + " years old"

print(Person("Gideon", 25)) # prints 'Gideon is 25 years old'

```

Listing 12.1: A simple Person class with a custom string representation

Classes that do not define their own string representation inherit the `__str__` attribute from their superclass(es). Every object has a string representation because all classes inherit `__str__` from `object`. The primitive `object.__str__` (specified in Rule A.1) uses an objects address to create a string representation that is unique for every object.

Evaluating the printed expression

First, the expression *e* to be printed is evaluated.

$$\Rightarrow \langle \Theta, \Gamma, S \mid \text{print}(e); \rangle$$

$$\Rightarrow \langle \Theta, \Gamma, S \mid \text{print}(o._str_()), e \rangle \quad (12.1)$$

Getting a string representation of the printed value

Next, we call the class function `__str__` of *a* to obtain a string representation of *a*. We don't need to check whether an object has a class function `__str__`, unlike for example the `__iter__` attribute in Rule 11.9.

$$\Rightarrow \langle \Theta, \Gamma, S \mid \text{print}(o._str_()), a \rangle$$

$$\Rightarrow \langle \Theta, \Gamma, S \mid \text{print}(o) \quad , a_f([a]) \rangle \quad (12.2)$$

where $a_f = \Phi_a^\Theta(_str_)$

Printing the string

Finally, the string representation $\Theta(a)$ of the object can be printed. Provided that the class function `__str__` returned a string.

$$\xrightarrow{\text{print } \Theta(a)} \langle \Theta, \Gamma, S | \text{print}(o), a \rangle \text{ if } a :_{\Theta} a_{\text{str}} \quad (12.3)$$

Faulty string representation

If the class function `__str__` did not return a string, we raise a type error.

$$\Rightarrow \langle \Theta, \Gamma, S | \text{print}(o), a \rangle \text{ if } \neg (a :_{\Theta} a_{\text{str}}) \quad (12.4)$$

13 Operators

The semantics of most operators are determined by the class function of their operands. For example the addition operator `+` is implemented by the class function `__add__` of the left operand. An expression $e_l + e_r$ could¹ be translated to $e_l.\text{__add__}([e_r])$.

Boolean operators are an exception to this rule. They are not implemented by class function but depend on the class function `__nonzero__` to coerce values to booleans.

13.1 Unary operators

Unary operators have only one operand. Evaluation takes three simple steps: evaluation of the operand, getting the operator's attribute from the operand's class, execute the class function.

Evaluating the operand of a unary operator expression

$$\begin{aligned} & \langle \Theta, \Gamma, S \quad , \circ e \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S | \circ \circ, e \rangle \end{aligned} \tag{13.1}$$

Evaluating the unary not operator

The `not` operator negates the value returned by the class function `__nonzero__` of its operand. If the operand has no attribute `__nonzero__`, we raise a `TypeError`.

$$\begin{aligned} & \langle \Theta, \Gamma, S | \text{not } \circ \quad , a \quad \rangle \\ \Rightarrow & \begin{cases} \langle \Theta, \Gamma, S | \text{not } \circ \text{__nonzero__} , a_f([a]) \rangle & \text{if } \text{__nonzero__} \in \Phi_a^\Theta \\ \langle \Theta, \Gamma, S \quad , \text{raise } a_{\text{TypeError}}() \rangle & \text{otherwise} \end{cases} \tag{13.2} \\ & \text{where } a_f = \Phi_a^\Theta(\text{__nonzero__}) \end{aligned}$$

¹This simple rewrite of the expression has a different order of evaluation and assumes `__getattr__` is not overridden.

Return the negated value of a not operand

If `__nonzero__` returned a boolean, we return its negated value. Otherwise, if `__nonzero__` returned something else, we raise a `TypeError`.

$$\begin{aligned} & \langle \Theta, \Gamma, S | \text{not } \circ_ \text{__nonzero_}, a \rangle \\ \Rightarrow & \begin{cases} \langle \Theta, \Gamma, S & , \neg \Theta(a) \rangle & \text{if } a : \Theta a_{\text{bool}} \\ \langle \Theta, \Gamma, S & , \text{raise } a_{\text{TypeError}}() \rangle & \text{otherwise} \end{cases} \end{aligned} \quad (13.3)$$

Executing the function of a unary operator

For other operators, we call the appropriate operator function.

$$\begin{aligned} & \langle \Theta, \Gamma, S | \bigcirc \circ, a \rangle \\ \Rightarrow & \begin{cases} \langle \Theta, \Gamma, S & , a_f([a]) \rangle & \text{if attribute of } \bigcirc \in \Phi_a^\Theta \\ \langle \Theta, \Gamma, S & , \text{raise } a_{\text{TypeError}}() \rangle & \text{otherwise} \end{cases} \\ & \text{where } a_f = \Phi_a^\Theta(\text{attribute of } \bigcirc) \end{aligned} \quad (13.4)$$

13.2 Binary operators

Most operators have two operands: a left and a right hand side operator. The semantics of the operator are defined by the left hand side operator.

Evaluating the left hand side of a binary operator

First, the left hand side operator is evaluated while the remaining expression is put on the stack.

$$\begin{aligned} & \langle \Theta, \Gamma, S & , e_l \bigcirc e_r \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S | \bigcirc e_r, e_l \rangle \end{aligned} \quad (13.5)$$

Evaluating the right hand side of a binary operator

The boolean operators `and` and `or` are lazy. They only evaluate their right hand side if needed. Therefore, we call `__nonzero__` for the `and` and `or` operators before `er` is evaluated. If `__nonzero__` is not defined, we raise a `TypeError`.

Other operators are executed after both their operators have been evaluated. Therefore, we evaluate e_r and store a_l on the stack.

$$\begin{aligned} & \langle \Theta, \Gamma, S | \circ \circ e_r, a_l \rangle \\ \Rightarrow & \begin{cases} \langle \Theta, \Gamma, S | \circ \circ_B e_r, a_f([a_l]) \rangle & \text{if } \circ \in [\text{and}, \text{or}] \wedge \text{_nonzero_} \in \Phi_{a_l}^\Theta \\ \langle \Theta, \Gamma, S, \text{raise } a_{\text{TypeError}}() \rangle & \text{if } \circ \in [\text{and}, \text{or}] \wedge \text{_nonzero_} \notin \Phi_{a_l}^\Theta \\ \langle \Theta, \Gamma, S | a_l \circ \circ, e_r \rangle & \text{otherwise} \end{cases} \quad (13.6) \\ & \text{where } a_f = \Phi_{a_l}^\Theta(\text{_nonzero_}) \end{aligned}$$

Evaluation of the boolean operators continues at Rules 13.9 and 13.10. The `in` operator continues at Rule 13.7. Other operators are evaluated by Rule 13.8.

Evaluating an `is` expression

The `is` operator determines whether its operands are the same object. This is achieved by comparing the addresses of the operands.

$$\begin{aligned} & \langle \Theta, \Gamma, S | a_l \text{is } \circ, a_r \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S, a_l \equiv a_r \rangle \end{aligned} \quad (13.7)$$

Calling the operator attribute

When both operands of the operator have been evaluated, the operator's class function of a_l is called with the operands as its arguments.

$$\begin{aligned} & \langle \Theta, \Gamma, S | a_l \circ \circ, a_r \rangle \\ \Rightarrow & \begin{cases} \langle \Theta, \Gamma, S, a_f([a_l, a_r]) \rangle & \text{if attribute of } \circ \in \Phi_{a_l}^\Theta \\ \langle \Theta, \Gamma, S, \text{raise } a_{\text{TypeError}}() \rangle & \text{otherwise} \end{cases} \quad (13.8) \\ & \text{where } a_f = \Phi_{a_l}^\Theta(\text{attribute of } \circ) \end{aligned}$$

Lazy execution of the `and` operator

The `and` expression returns e_r if the e_l is not zero (its `_nonzero_` returned `True`), otherwise it returns `False`. Note that the right hand side e_r can evaluate to any value.

$$\begin{aligned} & \langle \Theta, \Gamma, S | \circ \text{and}_B e_r, a \rangle \\ \Rightarrow & \begin{cases} \langle \Theta, \Gamma, S, e_r \rangle & \text{if } a : \Theta a_{\text{bool}} \wedge \Theta(a) \\ \langle \Theta, \Gamma, S, \text{False} \rangle & \text{if } a : \Theta a_{\text{bool}} \wedge \neg \Theta(a) \\ \langle \Theta, \Gamma, S, \text{raise } a_{\text{TypeError}}() \rangle & \text{otherwise} \end{cases} \quad (13.9) \end{aligned}$$

13 Operators

Lazy execution of the or operator

Conversely, the `or` expression returns e_r when e_l is zero and `True` otherwise.

$$\Rightarrow \left\langle \begin{array}{l} \langle \Theta, \Gamma, S | \circ_{\text{or}_B} e_r, a \rangle \\ \langle \Theta, \Gamma, S, \text{True} \rangle \text{ if } a : \Theta a_{\text{bool}} \wedge \Theta(a) \\ \langle \Theta, \Gamma, S, e_r \rangle \text{ if } a : \Theta a_{\text{bool}} \wedge \neg \Theta(a) \\ \langle \Theta, \Gamma, S, \text{raise } a_{\text{TypeError}}() \rangle \text{ otherwise} \end{array} \right\rangle \quad (13.10)$$

13.3 Container operators

Containers are values containing values, or actually addresses pointing to values. Individual items of a container can be retrieved, set, or deleted using the container operators. The semantics of those operators depend on the implementation of the container's class functions `__getitem__`, `__setitem__`, and `__delitem__`.

For example, lists, tuples and dictionaries are containers with their own implementation of the container attributes (see Appendix A). Not just built-in values are containers though, any object with the right class functions is considered a container.

Retrieving a container item

The subscription $e[e_i]$ retrieves an item with key or index e_i from a container e .

Evaluating the container expression

First, the container e is evaluated.

$$\Rightarrow \left\langle \begin{array}{l} \langle \Theta, \Gamma, S, e[e_i] \rangle \\ \langle \Theta, \Gamma, S | \circ [e_i], e \rangle \end{array} \right\rangle \quad (13.11)$$

Evaluating the key expression

Then, the index or key e_i is evaluated.

$$\Rightarrow \left\langle \begin{array}{l} \langle \Theta, \Gamma, S | \circ [e_i], a \rangle \\ \langle \Theta, \Gamma, S | a[\circ], e_i \rangle \end{array} \right\rangle \quad (13.12)$$

Calling `__getitem__` of a container

Finally, the container's class function `__getitem__` is called to retrieve the item.

$$\begin{aligned} & \langle \Theta, \Gamma, S | a[\circ], a_i \rangle \\ \Rightarrow & \begin{cases} \langle \Theta, \Gamma, S, a_f([a, a_i]) \rangle & \text{if } _getitem_ \in \Phi_a^\Theta \\ \langle \Theta, \Gamma, S, \text{raise } a_{\text{TypeError}}() \rangle & \text{otherwise} \end{cases} \\ & \text{where } a_f = \Phi_a^\Theta(_getitem_) \end{aligned} \quad (13.13)$$

Setting a container item

The subscription assignment statement $e[e_i] = e_v$ sets an item e_v in the container e at key or index e_i . Like the object attribute assignment (see Section 9.4) the expressions in the container assignment are not evaluated from left to right, but the value e_v is evaluated first.

Evaluating the right hand side expression

First, the item's value e_v is evaluated.

$$\begin{aligned} & \langle \Theta, \Gamma, S, e[e_i] = e_v; \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S | e[e_i] = \circ, e_v \rangle \end{aligned} \quad (13.14)$$

Evaluating the container expression

Then, the container e is evaluated.

$$\begin{aligned} & \langle \Theta, \Gamma, S | e[e_i] = \circ, a_v \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S | \circ [e_i] = a_v, e \rangle \end{aligned} \quad (13.15)$$

Evaluating the key expression

Next, the index or key e_i is evaluated.

$$\begin{aligned} & \langle \Theta, \Gamma, S | \circ [e_i] = a_v, a \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S | a[\circ] = a_v, e_i \rangle \end{aligned} \quad (13.16)$$

13 Operators

Calling `__setitem__`

Finally, the container's class function `__setitem__` is called to insert the new item.

$$\begin{aligned} & \langle \Theta, \Gamma, S | a[\circ] = a_v, a_i \rangle \\ \Rightarrow & \begin{cases} \langle \Theta, \Gamma, S & , a_f([a, a_i, a_v]) \rangle & \text{if } __setitem__ \in \Phi_a^\Theta \\ \langle \Theta, \Gamma, S & , \text{raise } a_{\text{TypeError}}() \rangle & \text{otherwise} \end{cases} & (13.17) \\ & \text{where } a_f = \Phi_a^\Theta(__setitem__) \end{aligned}$$

Deleting a container item

The subscription delete statement `del e[ei]` removes an item with key or index e_i from the container e .

Evaluating the container expression

First, the container e is evaluated.

$$\begin{aligned} & \langle \Theta, \Gamma, S & , \text{del } e[e_i]; \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S | \text{del } \circ [e_i], e \rangle & (13.18) \end{aligned}$$

Evaluating the key expression

Then, the key or index e_i is evaluated.

$$\begin{aligned} & \langle \Theta, \Gamma, S | \text{del } \circ [e_i], a \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S | \text{del } a[\circ] , e_i \rangle & (13.19) \end{aligned}$$

Calling `__delitem__`

Finally, the container's class attribute `__delitem__` is called to delete the item.

$$\begin{aligned} & \langle \Theta, \Gamma, S | \text{del } a[\circ], a_i \rangle \\ \Rightarrow & \begin{cases} \langle \Theta, \Gamma, S & , a_f([a, a_i]) \rangle & \text{if } __delitem__ \in \Phi_a^\Theta \\ \langle \Theta, \Gamma, S & , \text{raise } a_{\text{TypeError}}() \rangle & \text{otherwise} \end{cases} & (13.20) \\ & \text{where } a_f = \Phi_a^\Theta(__delitem__) \end{aligned}$$

14 Dynamic code execution

The `exec` and `import` statements are remarkably similar. Both dynamically parse and execute a block of code, and both statements can execute a block in an isolated environment. We will first explore the simpler `exec` statement, and then the `import` statement.

14.1 The `exec` statement

An `exec` statement `exec ep in eγ` executes a program `ep` in the environment `eγ`. The program `ep` should be a string value and the environment `eγ` a dictionary.

```
exec "v*2" in { "v" : 21 }
```

Listing 14.1: Executing an important question

Evaluating the string to be executed

First, the program string `ep` is evaluated.

$$\begin{aligned} & \langle \Theta, \Gamma, S, \text{exec } e_p \text{ in } e_\gamma; \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S | \text{exec } \circ \text{ in } e_\gamma, e_p \rangle \end{aligned} \tag{14.1}$$

Evaluating the execution environment

Then, the environment dictionary `eγ` is evaluated.

$$\begin{aligned} & \langle \Theta, \Gamma, S | \text{exec } \circ \text{ in } e_\gamma, a_p \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S | \text{exec } a_p \text{ in } \circ, e_\gamma \rangle \end{aligned} \tag{14.2}$$

Execute the string

1. If `ap` is a string and `γ` is a dictionary, we store the current environment on the stack, set the new environment and parse the program. The new environment is defined as the default environment `γ0` extended with the dictionary environment `γ`.
2. If either `ap` is not a block, or `γ` is not a dictionary, we raise a `TypeError`.

$$\begin{aligned} & \langle \Theta, \Gamma, S | \text{exec } a_p \text{ in } \circ, \gamma \rangle \\ \Rightarrow & \begin{cases} \langle \Theta, \gamma_0 | \gamma, S | \Gamma \vdash \text{exec } \circ, \text{parse } \Theta(a_p) \rangle & \text{if } a_p :_{\Theta} \Theta \wedge \gamma :_{\Theta} \Theta \\ \langle \Theta, \Gamma, S, \text{raise } a_{\text{TypeError}}() \rangle & \text{otherwise} \end{cases} \end{aligned} \quad (14.3)$$

The parse function will, if the string has the correct grammar, return a block. If the parser fails however, it will raise a `SyntaxError` exception.

Restoring the old environment

Once the program string has been executed, the old environment Γ' is reinstated and we return `None`.

$$\begin{aligned} & \langle \Theta, \Gamma, S | \Gamma' \vdash \text{exec } \circ, a \rangle \\ \Rightarrow & \langle \Theta, \Gamma', S, a_{\text{None}} \rangle \end{aligned} \quad (14.4)$$

Raising an exception in an execution environment

The old environment is also reinstated when an exception is raised during the execution of the program (see Chapter 10). Return, yield, break and continue statements however, can not escape the execution environment, because they cannot occur outside of their respective compound statements.

$$\begin{aligned} & \langle \Theta, \Gamma, S | \Gamma' \vdash \text{exec } \circ, \text{raise } a \rangle \\ \Rightarrow & \langle \Theta, \Gamma', S, \text{raise } a \rangle \end{aligned} \quad (14.5)$$

14.2 The import statement

Import statements initialize modules or packages, and include them in the current environment as module objects.

Packages are directories containing an initialization module `__init__.py` and usually a number of other modules. Modules are normal Python files that are implemented by executing them. The variables that have been bound after the initialization of a module, define the attributes of the new module object. Similarly, the initialized modules and packages in a package define the attributes of its module object.

Modules and packages are initialized only once. The first time a module or package is imported, it is stored in a dictionary value at $a_{\overline{m}}$. If a module is imported a second time, it will be retrieved from the dictionary $a_{\overline{m}}$. This mechanism also prevents the built-in modules from being redefined.

An import statement `import \bar{x}` imports a module or package \bar{x} . The dot separated list of names \bar{x} describes the path of a module or package, where the first names $x_1 \dots x_{n-1}$ are packages and the last name x_n is usually a module but can be a package.

When executing an import statement `import \bar{x}` with $|\bar{x}| > 1$, each package in the path is initialized before the full path is imported. It is as if we were to execute `import x_1` , `import $x_1.x_2$` , and so forth.

Executing an import statement

The import statement is executed by placing the import frame $\Gamma \vdash \text{import } \bar{x}$ on the stack. The import frame contains the environment Γ that is reinstated when the execution of the import statement has completed (or fails with an exception), and the path \bar{x} that will be imported.

$$\Rightarrow \left\langle \Theta, \Gamma, S \quad , \text{import } \bar{x}; \right\rangle \quad (14.6)$$

This first rule is always immediately followed by Rule 14.7.

Importing modules or packages

The import frame $\Gamma \vdash \text{import } \bar{x} \circ \bar{y}$ indicates that we have just imported the first part \bar{x} of the import statement's path, and have to import the remaining part \bar{y} . The topmost environment γ_p is the environment of the module that has just been imported. The address a , which is immediately discarded, is returned by the last statement of the previously imported module.

The environment γ_p , in which the module \bar{x} was executed, defines the module's attributes. This is not unlike the class definition statements (see Rule 9.1).

1. If a module or package x_m was already imported, its address is in list of previously imported modules \bar{m} . In this case we bind the module or package in the environment γ_p .
2. If x_m is a package (there is a directory by that name with an initialization module in it), we create a new environment a_m in which we execute the parsed initialization module. The environment a_m is bound to \bar{y} in the module γ_p and stored in the list of modules \bar{m} .
3. If we are importing the last element of the path and it is a module, we execute the parsed module in the environment a_m . The environment a_m is bound to \bar{y} in the module γ_p and stored in the list of modules \bar{m} .
4. Otherwise the import is invalid because a module or package could not be found, so we raise an `ImportError`.

$$\begin{aligned}
& \langle \Theta, \Gamma | \gamma_p, S | \Gamma \vdash \text{import } \bar{x}.\circ.(y : \bar{y}), a \rangle \\
\Rightarrow & \left\{ \begin{array}{l} \langle \Theta', \bar{m}(x), S | f \rangle, a_{\text{None}} \\ \langle \Theta'', \gamma_0 | a_m, S | f \rangle, \text{parse } p \\ \langle \Theta'', \gamma_0 | a_m, S | f \rangle, \text{parse } x.\text{py} \\ \langle \Theta, \Gamma, S \rangle, \text{raise } a_{\text{IE}}() \end{array} \right\} \begin{array}{l} \text{if } x \in \bar{m} \\ \text{if } x \notin \bar{m} \wedge x \text{ is package} \\ \text{if } x \notin \bar{m} \wedge x \text{ is module} \wedge |\bar{y}| \equiv 0 \\ \text{otherwise} \end{array} \\
\text{where } & a_m = \text{new address } \notin \Theta \\
& x = x_1 / \dots / x_n / y \\
& p = x / \text{__init__.py} \\
& \bar{m} = \Theta(a_{\bar{m}}) \\
& \Theta' = \Theta \oplus [\gamma_p \rightarrow \Theta(\gamma_p) \oplus [y \rightarrow \bar{m}(x)]] \\
& \Theta'' = \Theta \oplus [\gamma_p \rightarrow \Theta(\gamma_p) \oplus [y \rightarrow a_m] \\
& \quad, a_{\bar{m}} \rightarrow \bar{m} \oplus [x \rightarrow a_m] \\
& \quad, a_m \rightarrow \emptyset \\
& \quad] \\
& a_{\text{IE}} = a_{\text{ImportError}} \\
& f = \Gamma \vdash \text{import } (\bar{x} \# [y]).\circ.\bar{y}
\end{aligned} \tag{14.7}$$

Like the parser used for the `exec` statement (see Section 14.1), this parser returns a block when it succeeds and raises a `SyntaxError` exception when it fails. Note that we're parsing files here, not strings.

This rule applies repeatedly until either an exception is raised, or all the packages and the module have been imported. It is then followed by Rule 14.8.

Returning from the import statement

When all packages and the final module have been imported, we reinstate the environment γ and return `None`.

$$\begin{aligned}
& \langle \Theta, \Gamma | \gamma_p, S | \Gamma \vdash \text{import } \bar{x}.\circ.[], a \rangle \\
\Rightarrow & \langle \Theta, \Gamma, S, a_{\text{None}} \rangle
\end{aligned} \tag{14.8}$$

Raising an exception in an imported module

If an imported module raises an exception, the previous environment Γ' is reinstated and unwinding continues (see Chapter 10).

$$\begin{aligned}
& \langle \Theta, \Gamma, S | \Gamma' \vdash \text{import } \bar{x}.\circ.\bar{y}, \text{raise } a \rangle \\
\Rightarrow & \langle \Theta, \Gamma', S, \text{raise } a \rangle
\end{aligned} \tag{14.9}$$

15 Conclusion

We have created an executable operational semantics of a significant subset of Python. The subset, called *minpy*, covers all major constructs of Python. The semantics are easy to read, insofar the formal semantics of a serious programming language can be.

15.1 Literate programming

Developing an executable operational semantics was a challenging exercise of literate programming as introduced by Knuth [8].

In contrast to Knuth's WEB system, which transparently includes a program's source code into a document, we have tried to hide the Haskell source code in this document. Had we simply included the Haskell code, it would have been perceived as just an implementation of Python in Haskell, rather than an operational semantics of Python.

The Haskell source code was formatted with `lhs2TeX`. This system uses formatting rules to translate Haskell to \LaTeX , which can be edited independently from source code that is being formatted. For example, to reveal the Haskell code, one would simply remove the formatting rules and recompile the document.

We pushed `lhs2TeX` to its limits. It was designed for small formatting tweaks, not to completely transform the source code as we did. The sources had to be pre-processed by a simple Perl script and stick to strict coding conventions to get the current results.

The choice of Haskell as the specification language turned out rather well. Its declarative nature and almost mathematical syntax allowed us to write Haskell code that is close to the eventual formatted formulas.

In addition, Haskell provides many static checks which prevent many small errors. For example, undeclared variables or functions (often caused by typos) are detected, and the order and number of function arguments is verified by the type system.

The type system has also given us much more confidence in the correctness of the semantics. Despite its type inference, the type system of Haskell introduces some overhead. For example, in the semantics we rely on a convention to separate addresses from other types of values: they are all described by a variable a or γ . Haskell requires us to explicitly tag these variables, e.g. `Addr a`.

Since the type system also captures side effects, it prevents any accidental introduction of a hidden state that can affect the semantics. For example, the guard expressions used to express conditional rewrite rules cannot have side effects. After all, the act of searching for the rewrite rule that applies to a specific machine state should not influence the semantics.

Hiding Haskell was not always easy. Some concepts could have been described very concisely using a number of standard library functions, but to make this document self contained, many library functions could not be used. The formatting of Haskell also limits us to a subset of Haskell.

While we wanted to hide Haskell in the document, we've tried to keep the source code as

15 Conclusion

close to the formatted rules as possible. Haskell proved quite flexible in this respect. For example, the operator \oplus is the reformatted version of `<+>` in the Haskell code.

Because of the code conventions required for proper formatting and the restrictions on Haskell, the source code of the rewrite rules will not win any beauty contests (see Section 4). Nevertheless, it is almost certainly easier to read and write than any (non-executable) equivalent \LaTeX code.

Overall development on the semantics is quite a satisfying experience. When changing the semantics one would typically write some tests to explore the desired semantics, then implement a solution in Haskell, and finally, when the implementation passes the tests, ensure that it is properly formatted.

The Glasgow Haskell Compiler (GHC), directly supports literate Haskell. The source files do not need to be preprocessed by a separate tool but can be fed into GHC directly. Thus, one can disregard the restrictions while experimenting and worry about the formatting later.

15.2 The interpreter

The sources of this document and some supporting code that implements a parser, pretty printer, and an interactive interpreter, can be compiled by the Haskell compiler GHC to produce an interpreter. Like the interpreter of CPython, our interpreter has both a command-line modus and the ability to interpret files.

The following example shows a session of the interpreter in its command-line modus. First, we declare the factorial function in a single function definition statement. The statement is immediately executed if parsing succeeds. Next, we call the factorial function without an argument, which results in a `TypeError` exception. Finally we call it with a correct argument and exit the interpreter.

```
gideon@gideon-desktop:~$ minpy
>>> def fac(x) :
...     if x <= 0 :
...         return 1
...     else :
...         return fac(x-1) * x
...
>>> fac()
TypeError
>>> fac(4)
24
>>> exiting minpy
```

Without a standard library and a foreign function interface, the interpreter cannot be used for many serious applications. However, as a platform of experimentation and testing, the interpreter was invaluable.

The interpreter includes a tracer that has proven to be very helpful when debugging the rewrite rules. The tracer prints the changes of the abstract machine state.

The following interpreter session shows how the tracer is used. The tracer is first enabled by calling `enableTrace()` of the `sys` module. After a simple assignment statement the trace

shows how the value 1 is stored at address 1006, and the environment with address 22 is updated (in `state_5`) with a new mapping.

The interpreter always prints the result of the last statement if the result is not `None`. So when we execute the expression statement `a`, the value is retrieved (states 7 and 8) and printed (states 10 through 15).

```
>>> import sys
>>> sys.enableTrace()
state_1 = <heap_0, []:|:29:|:22, [], 0>
>>> a = 1
state_3 = <heap_2, envs_2, stack_2:|:AssFrame "a", 1>
state_4 = <heap_3<+>1006 |-> IntValue 1, envs_2, stack_3, 1006>
state_5 = <heap_4<+>22 |-> ModValue (fromList [("a",1006),("sys",25)]),
[]:|:29:|:22, stack_2, 0>
>>> a
state_7 = <heap_6, envs_6, stack_2, a>
state_8 = <heap_7, []:|:29:|:22, [], 1006>
state_10 = <heap_9, envs_9, stack_9:|:PrintFrame1, $1006>
state_11 = <heap_10, envs_9, stack_10, 1006>
state_12 = <heap_11, envs_9, stack_9:|:PrintFrame2, int.__str__:: (7)
-> ?([1006])>
state_13 = <heap_12, envs_9, stack_12, int.__str__([1006])>
state_14 = <heap_13<+>1007 |-> StringValue "1", envs_9, stack_12, 1007>
1
state_15 = <heap_14, envs_9, stack_9, 0>
```

Traces usually include a lot of addresses that are hard to keep track of. We implemented address expressions that allow us to handle addresses as first class object in the language for debugging. For example, if we had forgotten what the address 1006 pointed to we can retrieve it by entering the expression statement `$1006`.

```
>>> $1006
1
>>> exiting minpy
```

The interpreter does not only help when debugging but also provides insight in the semantics. Some aspects of the operational semantics ‘emerge’ from a number of rules. Judging where these emergent semantics occur and exploring their consequences, is often difficult. The interpreter allows us to explore these issues using examples.

For example, the interaction of the try-finally statement with the return statement (see Section 10.3) has some unexpected consequences for the semantics. It is quite hard to infer the semantics of this exceptional case from the rules. Running an example in the interpreter however, immediately reveals the emergent semantics.

Our interpreter runs considerably slower than CPython. This is no surprise. We paid little attention to the performance of the interpreter.

It should be noted though, that the difference is not a simple constant factor. The heap is implemented with a balanced binary tree, whereas CPython uses a direct mapping of memory

15 Conclusion

addresses to values. For a simple address lookup, the former has a complexity of $O(\log(n))$, whereas the latter has a constant complexity.

Due to the absence of any garbage collector in the interpreter our interpreter uses much more memory than CPython. Indeed, it will be hard to test any memory intensive programs and even a simple program such as `while True: pass` will eventually crash due to a lack of memory.

15.3 The test suite

As mentioned before, our executable operational semantics is testable: we can execute programs and verify whether the *minpy* interpreter gives the same answer as the CPython interpreter.

The CPython implementation is developed with a suite of regression tests. The regression tests could unfortunately not be used to test *minpy* because it relies on many Python features and libraries that are not available in *minpy*.

Fortunately the efforts to reverse engineer the semantics of CPython provided an excellent opportunity to write new tests. The resulting test suite consists of 134 tests and is completely written in *minpy*, including the test driver.

Beside testing the basic behaviour of constructs, the tests verify aspects such as control flow, order of evaluation, type/class requirements, and variable scoping. Each test covers a single aspect of a single construct.

We have been able to run the test suite for a number of implementations of Python (see Table 15.1). Testing other implementations improves confidence in the tested implementations.

Judging from the low number of failed tests across implementations, the other implementations seem to implement the same language, and conversely, *minpy* and the test-suite seem to capture a common subset of Python. Hopefully *minpy* can converge on the essential core of Python.

Implementations	Failed
CPython 2.5.2 (reference)	1
<i>minpy</i>	13
PyPy 1.0.0 build 41081	3
Iron Python 1.1.1 on .net 2.0.50727.42	15
Jython 2.2.1 on java1.6.0_0	17

Table 15.1: Number of failed tests for various Python implementations.

Since CPython is the reference implementation one would expect no errors. There is, however, one test that crashes the Python interpreter completely. The developers of CPython acknowledged that this was a bug and immediately solved the problem. None of the other implementations crashed in the same situation.

Our interpreter fails 13 tests, all of which are by definition errors in our semantics. Most of the errors are considered missing features. For example, lists in *minpy* raise an exception when a negative index is used. CPython on the other hand, treats negative indices as normal

indices on the reversed list with an offset of 1. Thus in CPython `[1, 2][-1]` evaluates to the value 2 but raises an exception in *minpy*.

Of course there are many more errors in *minpy* than have been tested or even discovered. Some of the known issues are listed in Section 15.4.

Interpreting the test results of other Python implementations is difficult. Some differences are intentional and can therefore not be considered bugs. Whether or not a failed test indicates a bug is a matter of opinion. Nevertheless we tried to interpret the results somewhat:

The PyPy interpreter failed three tests due to incompatibilities only; none of the failures indicated bugs. Iron Python failed quite a number of tests, but only five of those can be considered bugs. Of the failed tests on Jython, only three can be considered bugs.

15.4 Future work

There is an enormous body of work that can be based on this operational semantics. First of all, *minpy* can be expanded to cover all, or at least a more complete subset of Python.

The grammar of *minpy* and therefore the specification, is much simpler than Python's syntax (see Chapter 3). Most significantly, list and generator comprehensions, keyword arguments and other advanced function parameters, the with statement, and the global statement are missing.

Many optional elements and other syntactic sugar are also missing in *minpy*. This makes the semantics less verbose and allows us to concentrate on the most interesting aspects of Python. Nevertheless, for the sake of completeness, they should eventually be specified.

It should be noted that many constructs that could be considered syntactic sugar, actually have semantics that cannot completely be described by a simple transformation. For example, the augmented assignment $x += 1$ is not equivalent to $x = x + 1$: in the former case Python would call the class function `__iadd__`, but for the latter it would call `__setattr__` and `__add__`.

We are confident that all constructs of Python that are missing from *minpy* can be added without significant changes to the existing semantic rules. The structure of the abstract machine will not have to be changed, and the framework used to describe the rules will not need significant extensions.

Secondly, there are some errors in the current semantics that should be corrected. The test-suite pointed out some errors in the semantics of class initialization. Although we do use the class of a class –the metaclass– to construct classes, the mechanics of metaclasses have not been covered completely (see Section 9.1, specifically Rule 9.3).

Most of the built-in types have been specified to make the interpreter usable (see Chapter A). But the current semantics of *minpy*'s built-in types are not conform to Python. For example, *minpy* does not support coercion, and the dictionary class only accepts strings as keys.

Thirdly, the scope of the semantics (see Section 1.1) can be expanded to encompass more of Python's aspects:

- Garbage collection and its interaction with the operational semantics can be explored and formalised.
- Concurrent programming, notorious for its complexity, can be formalised.
- The interaction of Python with the outside world –the foreign function interface (FFI)– can be specified.

15 Conclusion

- The various reflective features can be formalised. It seems that most of these features can be specified fairly easily. For example, the `__dict__` attribute that exposes the attribute mappings of an object, would not require big changes.

For completeness one might consider to add so called old-style classes as well. It might be more productive to update the entire specification to a newer version of Python, that will does not support old style classes.

Fourthly, the interpreter, which will benefit from all improvements of the semantics, can be improved. Currently the error feedback is minimal. It would be very useful if raised exceptions also print a stack trace.

The performance of the interpreter can probably be improved, but it is doubtful it will ever achieve performance comparable to CPython without compromising the legibility of the semantics.

The interpreter could also be used as a basis for various other tools. The current basic tracing functionality could be extended to a fully fledged debugger for example.

Finally, our work opens many new avenues of research. The lack of a formal semantics for Python prevented any formal work on Python. Now however we can pursue various new directions of research:

- One can prove properties of Python with respect to our semantics. For example, we would very much like to prove that the execution of any possible Python program will never end in a state, other than the final state, that cannot be rewritten. Such a proof would have prevented the bug that we found in CPython.
- New tools for and analyses of Python programs can be developed more easily. For example, a tool that type-checks a Python program could be developed. Perhaps some properties of the type system could even be proven correct.
- Language extensions can be developed within our framework. Resulting in both a formal description of the extension and a proof of concept interpreter.

A Builtin values

A.1 The base objects `object` and `type`

```

 $o_{\text{object}} = \cdot$ 
  [ __getattr__ → object.__getattr__
    , __setattr__ → object.__setattr__
    , __delattr__ → object.__delattr__
    , __new__ → object.__new__
    , __init__ → object.__init__
    , __str__ → object.__str__
    , __bases__ →  $a_{()}()$ 
    , __class__ →  $a_{\text{type}}$ 
  ]

```

```

 $o_{\text{type}} = \cdot$ 
  [ __getattr__ → type.__getattr__
    , __bases__ →  $a_{(\text{object},)}$ 
    , __class__ →  $a_{\text{type}}$ 
    , __new__ → type.__new__
    , __call__ → type.__call__
    , __str__ → type.__str__
  ]

```

The default string representation

$$\Rightarrow \left\langle \Theta \oplus [a \rightarrow ("<\text{object at } " \text{ ++ } \textit{show } a_{\text{self}} \text{ ++ ">")}, \Gamma, S, a], \Gamma, S, a, \text{object}.__\text{str}__([a_{\text{self}}]) \right\rangle$$

where $a = \text{new address} \notin \Theta$

(A.1)

Default string representation for classes

$$\Rightarrow \left\langle \Theta \oplus [a \rightarrow ("<\text{class at } " \text{ ++ } \textit{show } a_{\text{self}} \text{ ++ ">")}, \Gamma, S, a], \Gamma, S, a, \text{type}.__\text{str}__([a_{\text{self}}]) \right\rangle$$

where $a = \text{new address} \notin \Theta$

(A.2)

A.2 The function and generator

```

 $O_{\text{generator}} = \cdot$ 
  [ __class__ →  $a_{\text{type}}$ 
    , __bases__ →  $a_{(\text{object},)}$ 
    , __next__ → generator.__next__
    , __send__ → generator.__send__
    , __iter__ → generator.__iter__
  ]

```

```

 $O_{\text{function}} = \cdot$ 
  [ __class__ →  $a_{\text{type}}$ 
    , __bases__ →  $a_{(\text{object},)}$ 
    , __call__ → function.__call__
  ]

```

A function as a callable object

$$\begin{aligned} & \langle \Theta, \Gamma, S, \text{function.}__\text{call}__([a]) \rangle \\ \Rightarrow & \langle \Theta, \Gamma, S, a \rangle \end{aligned} \tag{A.3}$$

The generator as an iterable

$$\begin{aligned} & \langle \Theta, \gamma, S, \text{generator.}__\text{iter}__([a_{\text{self}}]) \rangle \\ \Rightarrow & \langle \Theta, \gamma, S, a_{\text{self}} \rangle \end{aligned} \tag{A.4}$$

A.3 Integer values

The integer class

```

 $O_{\text{int}} = \cdot$ 
  [ __class__ →  $a_{\text{type}}$ 
    , __bases__ →  $a_{(\text{object},)}$ 
    , __new__ → int.__new__
    , __add__ → int.__add__: [ $a_{\text{int}}, a_{\text{int}}$ ]
    , __sub__ → int.__sub__: [ $a_{\text{int}}, a_{\text{int}}$ ]
    , __mul__ → int.__mul__: [ $a_{\text{int}}, a_{\text{int}}$ ]
    , __div__ → int.__div__: [ $a_{\text{int}}, a_{\text{int}}$ ]
    , __divmod__ → int.__divmod__: [ $a_{\text{int}}, a_{\text{int}}$ ]
    , __floordiv__ → int.__floordiv__: [ $a_{\text{int}}, a_{\text{int}}$ ]
    , __pow__ → int.__pow__: [ $a_{\text{int}}, a_{\text{int}}$ ]
  ]

```

```

, --neq--      → int.--neq--: [  $a_{\text{int}}$ ,  $a_{\text{int}}$  ]
, --eq--       → int.--eq--: [  $a_{\text{int}}$ ,  $a_{\text{int}}$  ]
, --lt--      → int.--lt--: [  $a_{\text{int}}$ ,  $a_{\text{int}}$  ]
, --le--      → int.--le--: [  $a_{\text{int}}$ ,  $a_{\text{int}}$  ]
, --gt--      → int.--gt--: [  $a_{\text{int}}$ ,  $a_{\text{int}}$  ]
, --ge--      → int.--ge--: [  $a_{\text{int}}$ ,  $a_{\text{int}}$  ]
, --str--     → int.--str--: [  $a_{\text{int}}$  ]
]

```

Literal integers

Evaluating an integer literal

$$\begin{aligned}
& \langle \Theta, \Gamma, S, i \rangle \\
\Rightarrow & \langle \Theta \oplus [a \rightarrow i], \Gamma, S, a \rangle \\
& \text{where } a = \text{new address} \notin \Theta
\end{aligned} \tag{A.5}$$

Integer primitives

The addition primitive

$$\begin{aligned}
& \langle \Theta, \Gamma, S, \text{int.--add--}([a_x, a_y]) \rangle \\
\Rightarrow & \langle \Theta, \Gamma, S, (\Theta(a_x) + \Theta(a_y)) \rangle
\end{aligned} \tag{A.6}$$

The equals primitive

$$\begin{aligned}
& \langle \Theta, \Gamma, S, \text{int.--eq--}([a_x, a_y]) \rangle \\
\Rightarrow & \langle \Theta, \Gamma, S, \Theta(a_x) \equiv \Theta(a_y) \rangle
\end{aligned} \tag{A.7}$$

The constructor primitive

$$\begin{aligned}
& \langle \Theta, \Gamma, S, \text{int.--new--}([a_c]) \rangle \\
\Rightarrow & \langle \Theta \oplus [a \rightarrow o], \Gamma, S, a \rangle \\
& \text{where } a = \text{new address} \notin \Theta \\
& \quad o = ([\text{--class--} \rightarrow a_c], 0)
\end{aligned} \tag{A.8}$$

The string representation

$$\begin{aligned}
& \langle \Theta, \Gamma, S, \text{int.--str--}([a_o]) \rangle \\
\Rightarrow & \langle \Theta \oplus [a \rightarrow (\text{show } \Theta(a_o))], \Gamma, S, a \rangle \\
& \text{where } a = \text{new address} \notin \Theta
\end{aligned} \tag{A.9}$$

A.4 Booleans

The boolean class

```
bool = ·
  [__class__  → atype
  , __bases__  → a(object,)
  , __nonzero__ → bool.__nonzero__: [abool]
  , __str__    → bool.__str__: [abool]
  ]
```

Literal booleans

Evaluating a boolean literal

$$\begin{aligned} & \langle \Theta, \gamma, S, b \rangle \\ \Rightarrow & \langle \Theta \oplus [a \rightarrow (b)], \gamma, S, a \rangle \\ & \text{where } a = \text{new address} \notin \Theta \end{aligned} \tag{A.10}$$

Boolean primitives

The boolean to boolean conversion primitive

$$\begin{aligned} & \langle \Theta, \gamma, S, \text{bool.__nonzero__}([a]) \rangle \\ \Rightarrow & \langle \Theta, \gamma, S, a \rangle \end{aligned} \tag{A.11}$$

The boolean to string conversion primitive

$$\begin{aligned} & \langle \Theta, \gamma, S, \text{bool.__str__}([a]) \rangle \\ \Rightarrow & \begin{cases} \langle \Theta, \gamma, S, \text{"True"} \rangle & \text{if } \Theta(a) \\ \langle \Theta, \gamma, S, \text{"False"} \rangle & \text{otherwise} \end{cases} \end{aligned} \tag{A.12}$$

A.5 Strings

The string class

```
Ostring = ·
  [__class__  → atype
  , __bases__  → a(object,)
  , __new__    → str.__new__
  , __add__    → str.__add__: [astr, astr]
  , __mul__    → str.__mul__: [astr, aint]
```

```

, __len__      → str.__len__: [a_str]
, __eq__      → str.__eq__: [a_str, a_str]
, __getitem__ → str.__getitem__: [a_str, a_int]
, __str__     → str.__str__: [a_str]
, "append"    → str.__add__: [a_str, a_str]
]

```

Literal strings

Evaluating a string literal

$$\begin{aligned}
& \langle \Theta, \gamma, S, t \rangle \\
\Rightarrow & \langle \Theta \oplus [a \rightarrow t], \gamma, S, a \rangle \\
\text{where } a = & \text{ new address } \notin \Theta
\end{aligned} \tag{A.13}$$

String primitives

The string constructor primitive

$$\begin{aligned}
& \langle \Theta, \gamma, S, \text{str.__new__}([a_c]) \rangle \\
\Rightarrow & \langle \Theta \oplus [a \rightarrow o], \gamma, S, a \rangle \\
\text{where } a = & \text{ new address } \notin \Theta \\
& o = ([_class_ \rightarrow a_c], [])
\end{aligned} \tag{A.14}$$

The string constructor primitive with starting string

$$\begin{aligned}
& \langle \Theta, \gamma, S, \text{str.__new__}([a_c, a_{\text{str}}]) \rangle \\
\Rightarrow & \langle \Theta \oplus [a \rightarrow o], \gamma, S, a \rangle \\
\text{where } a = & \text{ new address } \notin \Theta \\
& o = ([_class_ \rightarrow a_c], \Theta(a_{\text{str}}))
\end{aligned} \tag{A.15}$$

The length of string primitive

$$\begin{aligned}
& \langle \Theta, \gamma, S, \text{str.__len__}([a]) \rangle \\
\Rightarrow & \langle \Theta, \gamma, S, (|\Theta(a)|) \rangle
\end{aligned} \tag{A.16}$$

The string addition operator primitive

$$\begin{aligned}
& \langle \Theta, \gamma, S, \text{str.__add__}([a_{\text{self}}, a_{\text{other}}]) \rangle \\
\Rightarrow & \langle \Theta \oplus [a \rightarrow s], \gamma, S, a \rangle \\
\text{where } s = & (\Theta(a_{\text{self}}) \# \Theta(a_{\text{other}})) \\
& a = \text{ new address } \notin \Theta
\end{aligned} \tag{A.17}$$

A Builtin values

The string equals primitive

$$\begin{aligned} & \langle \Theta, \gamma, S, \text{str}.\text{eq}([a_x, a_y]) \rangle \\ \Rightarrow & \langle \Theta, \gamma, S, \Theta(a_x) \equiv \Theta(a_y) \rangle \end{aligned} \quad (\text{A.18})$$

The string get character at index primitive

$$\begin{aligned} & \langle \Theta, \gamma, S, \text{str}.\text{getitem}([a_{\text{self}}, a_i]) \rangle \\ \Rightarrow & \begin{cases} \langle \Theta \oplus [a \rightarrow [s !! i]], \gamma, S, a \rangle & \text{if } i < |s| \\ \langle \Theta, \gamma, S, \text{raise } a_{\text{IndexError}}() \rangle & \text{otherwise} \end{cases} \\ & \text{where } i = \Theta(a_i) \\ & \quad s = \Theta(a_{\text{self}}) \\ & \quad a = \text{new address} \notin \Theta \end{aligned} \quad (\text{A.19})$$

The string to string conversion primitive

$$\begin{aligned} & \langle \Theta, \gamma, S, \text{str}.\text{str}([a_o]) \rangle \\ \Rightarrow & \langle \Theta \oplus [a \rightarrow \Theta(a_o)], \gamma, S, a \rangle \\ & \text{where } a = \text{new address} \notin \Theta \end{aligned} \quad (\text{A.20})$$

A.6 Lists

The list class

```

olist = ·
  [__class__ → atype
  , __bases__ → a(object,)
  , __delitem__ → list.__delitem__: [alist, aint]
  , __getitem__ → list.__getitem__: [alist, aint]
  , __setitem__ → list.__setitem__: [alist, aint]
  , __new__ → list.__new__
  , __add__ → list.__add__
  , __len__ → list.__len__
  , __iter__ → list.__iter__
  , append → list.append
  ]

```

Literal strings

Evaluating a list literal

$$\begin{aligned} & \langle \Theta, \gamma, S, [\bar{e}] \rangle \\ \Rightarrow & \begin{cases} \langle \Theta, \gamma, S[[o, es'], e] \rangle & \text{if } \bar{e} = (e : es') \\ \langle \Theta \oplus [a \rightarrow []], \gamma, S, a \rangle & \text{otherwise} \end{cases} \\ & \text{where } a = \text{new address} \notin \Theta \end{aligned} \quad (\text{A.21})$$

Evaluating the elements of a list literal

$$\begin{aligned} & \langle \Theta, \gamma, S[[\bar{a}, o, \bar{e}], a] \rangle \\ \Rightarrow & \begin{cases} \langle \Theta, \gamma, S[[\bar{a} \# [a], o, es'], e] \rangle & \text{if } \bar{e} = (e : es') \\ \langle \Theta \oplus [a' \rightarrow [\bar{a} \# [a]]], \gamma, S, a' \rangle & \text{otherwise} \end{cases} \\ & \text{where } a' = \text{new address} \notin \Theta \end{aligned} \quad (\text{A.22})$$

List primitives**The list constructor primitive**

$$\begin{aligned} & \langle \Theta, \gamma, S, \text{list}.\text{new}([\bar{a}_c]) \rangle \\ \Rightarrow & \langle \Theta \oplus [a \rightarrow o], \gamma, S, a \rangle \\ & \text{where } a = \text{new address} \notin \Theta \\ & \quad o = ([\text{class} \rightarrow a_c], []) \end{aligned} \quad (\text{A.23})$$

The length of list primitive

$$\begin{aligned} & \langle \Theta, \gamma, S, \text{list}.\text{len}([a]) \rangle \\ \Rightarrow & \langle \Theta, \gamma, S, (|\Theta(a)|) \rangle \end{aligned} \quad (\text{A.24})$$

The get item at index primitive

$$\begin{aligned} & \langle \Theta, \gamma, S, \text{list}.\text{getitem}(a_{\text{self}}, a_i) \rangle \\ \Rightarrow & \begin{cases} \langle \Theta, \gamma, S, \bar{a} !! i \rangle & \text{if } i < |\bar{a}| \\ \langle \Theta, \gamma, S, \text{raise } a_{\text{IndexError}}() \rangle & \text{otherwise} \end{cases} \\ & \text{where } i = \Theta(a_i) \\ & \quad \bar{a} = \Theta(a_{\text{self}}) \end{aligned} \quad (\text{A.25})$$

The set item at index primitive

$$\begin{aligned} & \langle \Theta, \gamma, S, \text{list}.\text{setitem}(a_{\text{self}}, a_i, a) \rangle \\ \Rightarrow & \begin{cases} \langle \Theta \oplus [a_{\text{self}} \rightarrow \bar{a}'], \gamma, S, a_{\text{None}} \rangle & \text{if } i < |\bar{a}| \\ \langle \Theta, \gamma, S, \text{raise } a_{\text{IndexError}}() \rangle & \text{otherwise} \end{cases} \\ & \text{where } i = \Theta(a_i) \\ & \quad \bar{a} = \Theta(a_{\text{self}}) \\ & \quad \bar{a}' = [\dots, a_{i-1}, a, a_{i+1}, \dots] \end{aligned} \quad (\text{A.26})$$

The delete item at index primitive

$$\begin{aligned}
 & \langle \Theta \quad \quad \quad , \gamma, S, \text{list}.\text{delitem}([a_{\text{self}}, a_i]) \rangle \\
 \Rightarrow & \left\{ \begin{array}{l} \langle \Theta \oplus [a_{\text{self}} \rightarrow \bar{a}'], \gamma, S, a_{\text{None}} \\ \langle \Theta \quad \quad \quad , \gamma, S, \text{raise } a_{\text{IndexError}}() \end{array} \right\} \begin{array}{l} \text{if } i < |\bar{a}| \\ \text{otherwise} \end{array} \\
 \text{where } & i = \Theta(a_i) \\
 & \bar{a} = \Theta(a_{\text{self}}) \\
 & \bar{a}' = [\dots, a_{i-1}, a_{i+1}, \dots]
 \end{aligned} \tag{A.27}$$

The append to list primitive

$$\begin{aligned}
 & \langle \Theta \quad \quad \quad , \gamma, S, \text{list}.\text{append}([a_{\text{self}}, a]) \rangle \\
 \Rightarrow & \langle \Theta \oplus [a_{\text{self}} \rightarrow \bar{a}], \gamma, S, a_{\text{None}} \rangle \\
 \text{where } & \bar{a} = [\Theta(a_{\text{self}}) \# [a]]
 \end{aligned} \tag{A.28}$$

The list addition operator primitive

$$\begin{aligned}
 & \langle \Theta \quad \quad \quad , \gamma, S, \text{list}.\text{add}([a_{\text{self}}, a_{\text{other}}]) \rangle \\
 \Rightarrow & \langle \Theta \oplus [a \rightarrow \bar{a}], \gamma, S, a \rangle \\
 \text{where } & \bar{a} = [\Theta(a_{\text{self}}) \# \Theta(a_{\text{other}})] \\
 & a = \text{new address} \notin \Theta
 \end{aligned} \tag{A.29}$$

A.7 Tuples

A.8 The tuple class

```

otuple = .
  [__class__  → atype
  , __bases__ → a(object,)
  , __getitem__ → tuple.__getitem__: [atuple, aint]
  , __new__    → tuple.__new__
  , __len__    → tuple.__len__
  ]

```

Literal tuples

Evaluating a tuple literal

$$\begin{aligned}
 & \langle \Theta \quad \quad \quad , \gamma, S \quad \quad \quad , (\bar{e}) \rangle \\
 \Rightarrow & \left\{ \begin{array}{l} \langle \Theta \quad \quad \quad , \gamma, S | (\circ, es'), e \rangle \\ \langle \Theta \oplus [a \rightarrow (,)], \gamma, S \quad \quad \quad , a \rangle \end{array} \right\} \begin{array}{l} \text{if } \bar{e} = (e : es') \\ \text{otherwise} \end{array} \\
 \text{where } & a = \text{new address} \notin \Theta
 \end{aligned} \tag{A.30}$$

Evaluating the elements of a tuple literal

$$\begin{aligned}
& \langle \Theta, \gamma, S | (\bar{a}, o, \bar{e}), a \rangle \\
\Rightarrow & \begin{cases} \langle \Theta, \gamma, S | (\bar{a} + [a], o, \bar{e}'), e \rangle & \text{if } \bar{e} = (e : \bar{e}') \\ \langle \Theta \oplus [a' \rightarrow (\bar{a} + [a])], \gamma, S, a' \rangle & \text{otherwise} \end{cases} \\
& \text{where } a' = \text{new address} \notin \Theta
\end{aligned} \tag{A.31}$$

Tuple primitives**The tuple constructor primitive**

$$\begin{aligned}
& \langle \Theta, \gamma, S, \text{tuple}.\text{new}([a_c]) \rangle \\
\Rightarrow & \langle \Theta \oplus [a \rightarrow o], \gamma, S, a \rangle \\
& \text{where } a = \text{new address} \notin \Theta \\
& \quad o = ([\text{__class__} \rightarrow a_c], ([]))
\end{aligned} \tag{A.32}$$

The length of tuple primitive

$$\begin{aligned}
& \langle \Theta, \gamma, S, \text{tuple}.\text{len}([a]) \rangle \\
\Rightarrow & \langle \Theta, \gamma, S, (|\Theta(a)|) \rangle
\end{aligned} \tag{A.33}$$

The get item at index primitive

$$\begin{aligned}
& \langle \Theta, \gamma, S, \text{tuple}.\text{getitem}([a_{\text{self}}, a_i]) \rangle \\
\Rightarrow & \begin{cases} \langle \Theta, \gamma, S, \bar{a} !! i \rangle & \text{if } i < |\bar{a}| \\ \langle \Theta, \gamma, S, \text{raise } a_{\text{IndexError}} \rangle & \text{otherwise} \end{cases} \\
& \text{where } i = \Theta(a_i) \\
& \quad \bar{a} = \Theta(a_{\text{self}})
\end{aligned} \tag{A.34}$$

A.9 Dictionaries**Dictionary class**

```

odict = .
  [__class__  → atype
  , __bases__ → a(object,)
  , __getitem__ → dict.__getitem__: [adict, astr]
  , __setitem__ → dict.__setitem__: [adict, astr, aobject]
  , __delitem__ → dict.__delitem__: [adict, astr]
  , __new__    → dict.__new__
  ]

```

Dictionary literals

Evaluating a dictionary literal

$$\begin{aligned} & \langle \Theta, \gamma, S, \bar{e} \rangle \\ \Rightarrow & \begin{cases} \langle \Theta, \gamma, S | \{\emptyset, \circ : \bar{e}_{1;2}, \text{tail } \bar{e}\}, \bar{e}_{1;1} \rangle & \text{if } |\bar{e}| > 0 \\ \langle \Theta \oplus [a \rightarrow \emptyset], \gamma, S, a \rangle & \text{otherwise} \end{cases} \\ & \text{where } a = \text{new address} \notin \Theta \end{aligned} \quad (\text{A.35})$$

Evaluating the value of a key-value pair

$$\begin{aligned} & \langle \Theta, \gamma, S | \{d, \circ : e_v, \bar{e}\}, a_k \rangle \\ \Rightarrow & \langle \Theta, \gamma, S | \{d, a_k : \circ, \bar{e}\}, e_v \rangle \end{aligned} \quad (\text{A.36})$$

Evaluating key-value pairs in a dictionary literal

$$\begin{aligned} & \langle \Theta, \gamma, S | \{d, a_k : \circ, \bar{e}\}, a_v \rangle \\ \Rightarrow & \begin{cases} \langle \Theta, \gamma, S | \{d \oplus [(k \rightarrow a_v)], \circ : \bar{e}_{1;2}, \text{tail } \bar{e}\}, \bar{e}_{1;1} \rangle & \text{if } |\bar{e}| > 0 \\ \langle \Theta \oplus [a \rightarrow d \oplus [k \rightarrow a_v]], \gamma, S, a \rangle & \text{otherwise} \end{cases} \\ & \text{where } a = \text{new address} \notin \Theta \\ & \quad k = \Theta(a_k) \end{aligned} \quad (\text{A.37})$$

Dictionary primitives

The dictionary constructor primitive

$$\begin{aligned} & \langle \Theta, \gamma, S, \text{dict}._\text{new}_-\langle [a_c] \rangle \rangle \\ \Rightarrow & \langle \Theta \oplus [a \rightarrow ([_class_ \rightarrow a_c], \emptyset)], \gamma, S, a \rangle \\ & \text{where } a = \text{new address} \notin \Theta \end{aligned} \quad (\text{A.38})$$

The dictionary constructor primitive with starting dictionary

$$\begin{aligned} & \langle \Theta, \gamma, S, \text{dict}._\text{new}_-\langle [a_c, a_d] \rangle \rangle \\ \Rightarrow & \langle \Theta \oplus [a \rightarrow ([_class_ \rightarrow a_c], \Theta(a_d))], \gamma, S, a \rangle \\ & \text{where } a = \text{new address} \notin \Theta \end{aligned} \quad (\text{A.39})$$

The set item with key primitive

$$\begin{aligned} & \langle \Theta, \gamma, S, \text{dict}._\text{setitem}_-\langle [a_{\text{self}}, a_i, a_v] \rangle \rangle \\ \Rightarrow & \langle \Theta \oplus [a_{\text{self}} \rightarrow d'], \gamma, S, a_{\text{None}} \rangle \\ & \text{where } k = \Theta(a_i) \\ & \quad d' = (\Omega_{a_{\text{self}}}^\Theta, \Theta(a_{\text{self}}) \oplus [k \rightarrow a_v]) \end{aligned} \quad (\text{A.40})$$

The get item with key primitive

$$\begin{aligned}
& \langle \Theta, \gamma, S, \text{dict}._\text{getitem}__([a_{\text{self}}, a_i]) \rangle \\
\Rightarrow & \left\{ \begin{array}{l} \langle \Theta, \gamma, S, \Theta(a_{\text{self}})(i) \rangle \\ \langle \Theta, \gamma, S, \text{raise } a_{\text{IndexError}} \rangle \end{array} \right. \begin{array}{l} \text{if } i \in \Theta(a_{\text{self}}) \\ \text{otherwise} \end{array} \\
& \text{where } i = \Theta(a_i)
\end{aligned} \tag{A.41}$$

The delete item with key primitive

$$\begin{aligned}
& \langle \Theta, \gamma, S, \text{dict}._\text{delitem}__([a_{\text{self}}, a_i]) \rangle \\
\Rightarrow & \langle \Theta \oplus [a_{\text{self}} \rightarrow d'], \gamma, S, a_{\text{None}} \rangle \\
& \text{where } k = \Theta(a_i) \\
& \quad d' = (\Omega_{a_{\text{self}}}^\Theta, \Theta(a_{\text{self}}) \ominus k)
\end{aligned} \tag{A.42}$$

A Builtin values

Rules

6.1	Executing a non-empty block statement	23
6.2	Executing an empty block	23
6.3	Continuing execution of a block	23
6.4	Pass statements	24
6.5	Expression statements	24
7.1	Executing an assignment statement	25
7.2	Binding a variable	25
7.3	Evaluating a variable expression	26
7.4	Executing a variable deletion	26
8.1	Executing a function definition statement	28
8.2	Evaluating the function expression	29
8.3	Evaluating the first argument expression	29
8.4	Evaluating the remaining argument expressions	29
8.5	Type checking function arguments	30
8.6	Execution of a function	30
8.7	Returning from the end of a function	31
8.8	Executing a return statement	32
8.9	Start stack unwinding	32
8.10	Unwinding the stack	32
8.11	Getting the next yielded value	33
8.12	Getting the next value with yield value	33
8.13	Evaluating a yield expression	34
8.14	Start unwinding of the stack	34
8.15	Unwinding the stack	34
8.16	Stopping at the end of a generator	34
9.1	Evaluating the bases of a new class	38
9.2	Executing the body of a class definition	38
9.3	Creating a new class	38
9.4	The default class constructor	39
9.5	Initializing the new class	39
9.6	The default constructor	40
9.7	The default object creation	40
9.8	From creation to initialization	40
9.9	The default object initialization	40
9.10	Evaluating the attribute access expression	41
9.11	Calling <code>__getattr__</code>	41
9.12	Getting attributes and instantiating function attributes	41
9.13	Executing an instantiated function	42
9.14	Getting class members and inserting type checks	42
9.15	Evaluating the right hand side of an attribute assignment	43

A Builtin values

9.16	Evaluating the left hand side of an attribute assignment	43
9.17	Calling <code>__setattr__</code>	43
9.18	The default implementation of <code>__setattr__</code> for objects	43
9.19	Evaluating the left hand side of an attribute deletion	44
9.20	Calling <code>__delattr__</code>	44
9.21	The default implementation of <code>__delattr__</code>	44
10.1	Executing the raise statement	45
10.2	Start unwinding of the stack	46
10.3	Raising an exception	46
10.4	Raising an exception in a function	46
10.5	Raising an exception in a generator	46
10.6	Executing the try-except statement	47
10.7	Popping an exception clause while unwinding	47
10.8	Catching an exception	47
10.9	Popping the exception marker	48
10.10	Executing the try-finally statement	48
10.11	Popping a finally clause while unwinding	48
10.12	Popping the finally marker	49
11.1	Evaluating the test expression of an if statement	51
11.2	Testing the test value of an if statement	51
11.3	Executing the body of an if statement	52
11.4	Evaluating the test expression of a while statement	52
11.5	Testing the test value of a while statement	53
11.6	Executing the body of a while statement	53
11.7	Repeating a while statement	53
11.8	Executing the for statement	54
11.9	Get the iterator of the iterable	54
11.10	Start the for loop with the iterator	54
11.11	Starting an iteration of the for loop	55
11.12	Executing the body of a for loop	55
11.13	Stopping the for loop	55
11.14	Rewinding the stack for a break statement	55
11.15	Rewinding the stack for a continue statement	56
12.1	Evaluating the printed expression	57
12.2	Getting a string representation of the printed value	57
12.3	Printing the string	58
12.4	Faulty string representation	58
13.1	Evaluating the operand of a unary operator expression	59
13.2	Evaluating the unary not operator	59
13.3	Return the negated value of a not operand	60
13.4	Executing the function of a unary operator	60
13.5	Evaluating the left hand side of a binary operator	60
13.6	Evaluating the right hand side of a binary operator	60
13.7	Evaluating an <i>is</i> expression	61
13.8	Calling the operator attribute	61
13.9	Lazy execution of the and operator	61
13.10	Lazy execution of the or operator	62

13.11	Evaluating the container expression	62
13.12	Evaluating the key expression	62
13.13	Calling <code>__getitem__</code> of a container	63
13.14	Evaluating the right hand side expression	63
13.15	Evaluating the container expression	63
13.16	Evaluating the key expression	63
13.17	Calling <code>__setitem__</code>	64
13.18	Evaluating the container expression	64
13.19	Evaluating the key expression	64
13.20	Calling <code>__delitem__</code>	64
14.1	Evaluating the string to be executed	65
14.2	Evaluating the execution environment	65
14.3	Execute the string	65
14.4	Restoring the old environment	66
14.5	Raising an exception in an execution environment	66
14.6	Executing an import statement	67
14.7	Importing modules or packages	67
14.8	Returning from the import statement	68
14.9	Raising an exception in an imported module	68
A.1	The default string representation	75
A.2	Default string representation for classes	75
A.3	A function as a callable object	76
A.4	The generator as an iterable	76
A.5	Evaluating an integer literal	77
A.6	The addition primitive	77
A.7	The equals primitive	77
A.8	The constructor primitive	77
A.9	The string representation	77
A.10	Evaluating a boolean literal	78
A.11	The boolean to boolean conversion primitive	78
A.12	The boolean to string conversion primitive	78
A.13	Evaluating a string literal	79
A.14	The string constructor primitive	79
A.15	The string constructor primitive with starting string	79
A.16	The length of string primitive	79
A.17	The string addition operator primitive	79
A.18	The string equals primitive	80
A.19	The string get character at index primitive	80
A.20	The string to string conversion primitive	80
A.21	Evaluating a list literal	80
A.22	Evaluating the elements of a list literal	81
A.23	The list constructor primitive	81
A.24	The length of list primitive	81
A.25	The get item at index primitive	81
A.26	The set item at index primitive	81
A.27	The delete item at index primitive	82
A.28	The append to list primitive	82

A Builtin values

A.29	The list addition operator primitive	82
A.30	Evaluating a tuple literal	82
A.31	Evaluating the elements of a tuple literal	83
A.32	The tuple constructor primitive	83
A.33	The length of tuple primitive	83
A.34	The get item at index primitive	83
A.35	Evaluating a dictionary literal	84
A.36	Evaluating the value of a key-value pair	84
A.37	Evaluating key-value pairs in a dictionary literal	84
A.38	The dictionary constructor primitive	84
A.39	The dictionary constructor primitive with starting dictionary	84
A.40	The set item with key primitive	84
A.41	The get item with key primitive	85
A.42	The delete item with key primitive	85

Bibliography

- [1] D. Ancona, M. Ancona, A. Cuni, and N.D. Matsakis. RPython: a step towards reconciling dynamically and statically typed OO languages. *Proceedings of the 2007 symposium on Dynamic languages*, pages 53–64, 2007.
- [2] K. Barrett, B. Cassels, P. Haahr, D.A. Moon, K. Playford, and P.T. Withington. A monotonic superclass linearization for Dylan. *Proceedings of the 11th ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications*, pages 69–82, 1996.
- [3] The CPython interpreter. <http://www.python.org>.
- [4] New-style Classes. [ref:http://www.python.org/doc/newstyle/](http://www.python.org/doc/newstyle/).
- [5] P. Hudak, J. Hughes, S.P. Jones, and P. Wadler. A history of Haskell: being lazy with class. In *Proceedings of the third ACM SIGPLAN conference on History of programming languages*, pages 12–1. ACM Press New York, NY, USA, 2007.
- [6] The IronPython interpreter. <http://www.codeplex.com/Wiki/View.aspx?ProjectName=IronPython>.
- [7] The Jython interpreter. <http://www.jython.org>.
- [8] DE Knuth. Literate Programming. *The Computer Journal*, 27(2):97–111, 1984.
- [9] Andres Löh. lhs2TeX a literal programming tool for Haskell. <http://people.cs.uu.nl/andres/lhs2tex/>.
- [10] G.D. Plotkin. A structural approach to operational semantics. Report DAIMI FN-19. *Computer Science Department, Aarhus University*, 1981.
- [11] The PyPy interpreter. <http://pypy.org>.
- [12] Michele Simionato. The Python 2.3 method resolution order. <http://www.python.org/download/releases/2.3/mro/>.
- [13] G. van Rossum and F.L. Drake Jr. *Python Reference Manual*. Centrum voor Wiskunde en Informatica, 1995.
- [14] Guido van Rossum. Overview of the Python enhancement proposals. <http://www.python.org/dev/peps/>.
- [15] Guido van Rossum. Python Reference Manual. <http://docs.python.org/ref/>.