

**How To
Create
Your Own
Freaking
Awesome
Programming
Language**

TABLE OF CONTENT

1. [Table of Content](#)
2. [Introduction](#)
 1. [Summary](#)
 2. [About The Author](#)
 3. [Before We Begin](#)
3. [Overview](#)
 1. [The Four Parts of a Language](#)
 2. [Meet Awesome: Our Toy Language](#)
4. [Lexer](#)
 1. [Lex \(Flex\)](#)
 2. [Ragel](#)
 3. [Python Style Indentation For Awesome](#)
 4. [Do It Yourself I](#)
5. [Parser](#)
 1. [Bison \(Yacc\)](#)
 2. [Lemon](#)
 3. [ANTLR](#)
 4. [PEGs](#)
 5. [Operator Precedence](#)
 6. [Connecting The Lexer and Parser in Awesome](#)
 7. [Do It Yourself II](#)
6. [Runtime Model](#)
 1. [Procedural](#)
 2. [Class-based](#)
 3. [Prototype-based](#)
 4. [Functional](#)
 5. [Our Awesome Runtime](#)
 6. [Do It Yourself III](#)

7. [Interpreter](#)
 1. [Do It Yourself IV](#)
8. [Compilation](#)
 1. [Using LLVM from Ruby](#)
 2. [Compiling Awesome to Machine Code](#)
9. [Virtual Machine](#)
 1. [Byte-code](#)
 2. [Types of VM](#)
 3. [Prototyping a VM in Ruby](#)
10. [Going Further](#)
 1. [Homoiconicity](#)
 2. [Self-Hosting](#)
 3. [What's Missing?](#)
11. [Resources](#)
 1. [Books & Papers](#)
 2. [Events](#)
 3. [Forums and Blogs](#)
 4. [Interesting Languages](#)
12. [Solutions to Do It Yourself](#)
 1. [Solutions to Do It Yourself I](#)
 2. [Solutions to Do It Yourself II](#)
 3. [Solutions to Do It Yourself III](#)
 4. [Solutions to Do It Yourself IV](#)
13. [Appendix: Mio, a minimalist homoiconic language](#)
 1. [Homoicowhat?](#)
 2. [Messages all the way down](#)
 3. [The Runtime](#)
 4. [Implementing Mio in Mio](#)
 5. [But it's ugly](#)
14. [Farewell!](#)

Published November 2011.

Cover background image © [Asja Boros](#)

Content of this book is © Marc-André Cournoyer. All right reserved. This eBook copy is for a single user. You may not share it in any way unless you have written permission of the author.

INTRODUCTION

When you don't create things, you become defined by your tastes rather than ability. Your tastes only narrow & exclude people. So create.

- *Why the Lucky Stiff*

Creating a programming language is the perfect mix of art and science. You're creating a way to express yourself, but at the same time applying computer science principles to implement it. Since we aren't the first ones to create a programming language, some well established tools are around to ease most of the exercise. Nevertheless, it can still be hard to create a fully functional language because it's impossible to predict all the ways in which someone will use it. That's why making your own language is such a great experience. You never know what someone else might create with it!

I've written this book to help other developers discover the joy of creating a programming language. Coding my first language was one of the most amazing experiences in my programming career. I hope you'll enjoy reading this book, but mostly, I hope you'll write your own programming language.

If you find an error or have a comment or suggestion while reading the following pages, please send me an email at macournoyer@gmail.com.

SUMMARY

This book is divided into ten sections that will walk you through each step of language-building. Each section will introduce a new concept and then apply its principles to a language that we'll build together throughout the book. All technical chapters end with a *Do It Yourself* section that suggest some language-extending exercises. You'll find solutions to those at the end of this book.

Our language will be dynamic and very similar to Ruby and Python. All of the code will be in Ruby, but I've put lots of attention to keep the code as simple as possible so that you can understand what's happening even if you don't know Ruby.

The focus of this book is not on how to build a production-ready language. Instead, it should serve as an introduction in building your first toy language.

ABOUT THE AUTHOR

I'm Marc-André Cournoyer, a coder from Montréal, Québec passionate about programming languages and tequila, but usually not at the same time.

I coded [tinyrb](#), the smallest Ruby Virtual Machine, [Min](#), a toy language running on the JVM, [Thin](#), the high performance Ruby web server, and a bunch of other stuff. You can find most of my projects on [GitHub](#).

You can find me online, [blogging](#) or [tweeting](#) and offline, snowboarding and learning guitar.

BEFORE WE BEGIN

You should have received a `code.zip` file with this book including all the code samples. To run the examples you must have the following installed:

- [Ruby 1.8.7 or 1.9.2](#)
- Racc 1.4.6, install with: `gem install racc -v=1.4.6` (optional, to recompile the parser in the exercises).

Other versions might work, but the code was tested with those.

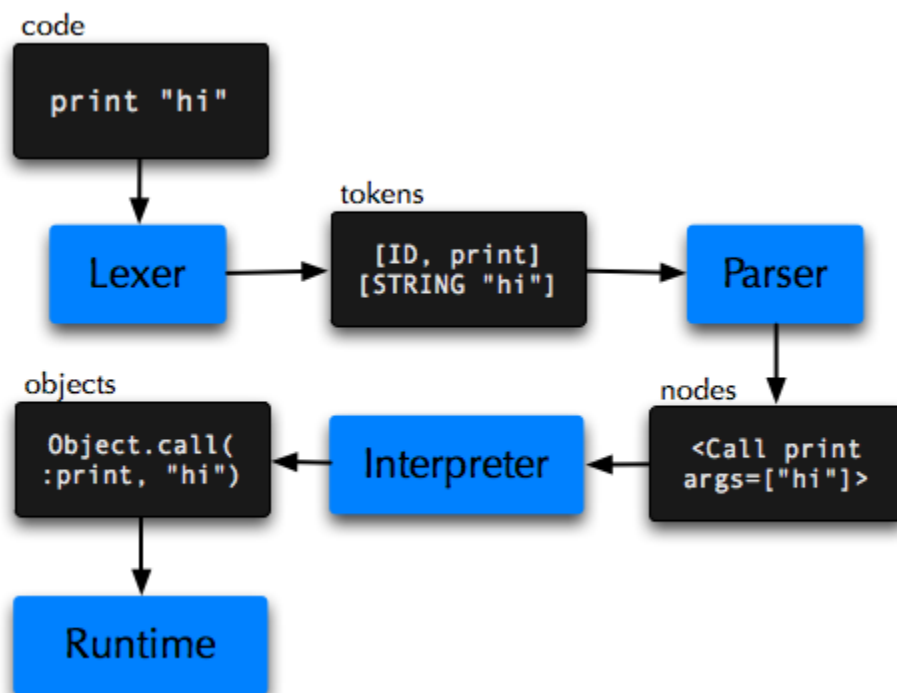
OVERVIEW

Although the way we design a language has evolved since its debuts, most of the core principles haven't changed. Contrary to Web Applications, where we've seen a growing number of frameworks, languages are still built from a lexer, a parser and a compiler. Some of the best books in this sphere have been written a long time ago: [Principles of Compiler Design](#) was published in 1977 and [Smalltalk-80: The Language and its Implementation](#) in 1983.

THE FOUR PARTS OF A LANGUAGE

Most dynamic languages are designed in four parts that work in sequence: the lexer, the parser, the interpreter and the runtime. Each one transforms the input of its predecessor until the code is run. Figure 1 shows an overview of this process. A chapter of this book is dedicated to each part.

Figure 1



MEET AWESOME: OUR TOY LANGUAGE

The language we'll be coding in this book is called Awesome, because it is!

It's a mix of Ruby syntax and Python's indentation:

```
1 class Awesome:
2     def name:
3         "I'm Awesome"
4
5     def awesomeness:
6         100
7
8 awesome = Awesome.new
9 print(awesome.name)
10 print(awesome.awesomeness)
```

A couple of rules for our language:

- As in Python, blocks of code are delimited by their indentation.
- Classes are declared with the `class` keyword.
- Methods can be defined anywhere using the `def` keyword.
- Identifiers starting with a capital letter are constants which are globally accessible.
- Lower-case identifiers are local variables or method names.
- If a method takes no arguments, parenthesis can be skipped, much like in Ruby.
- The last value evaluated in a method is its return value.
- Everything is an object.

Some parts will be incomplete, but the goal with a toy language is to educate, experiment and set the foundation on which to build more.

LEXER

The lexer, or scanner, or tokenizer is the part of a language that converts the input, the code you want to execute, into tokens the parser can understand.

Let's say you have the following code:

```
1 print "I ate",  
2     3,  
3     pies
```

Once this code goes through the lexer, it will look something like this:

```
1 [IDENTIFIER print] [STRING "I ate"] [COMMA]  
2         [NUMBER 3] [COMMA]  
3         [IDENTIFIER pies]
```

What the lexer does is split the code and tag each part with the type of token it contains. This makes it easier for the parser to operate since it doesn't have to bother with details such as parsing a floating point number or parsing a complex string with escape sequences (`\n`, `\t`, etc.).

Lexers can be implemented using regular expressions, but more appropriate tools exists.

LEX (FLEX)

[Flex](#) is a modern version of [Lex](#) (that was coded by Eric Schmidt, CEO of Google, by the way) for generating C lexers. Along with Yacc, Lex is the most commonly used lexer for parsing.

It has been ported to several target languages.

- [Rex for Ruby](#)
- [JFlex for Java](#)

Lex and friends are not lexers per se. They are lexers compilers. You supply it a grammar and it will output a lexer. Here's what that grammar looks like:

```
1 BLANK      [ \t\n]+
2
3 %%
4
5 // Whitespace
6 {BLANK}    /* ignore */
7
8 // Literals
9 [0-9]+     yylval = atoi(yytext); return T_NUMBER;
10
11 // Keywords
12 "end"      yyval = yytext; return T_END;
13 // ...
```

On the left side a regular expression defines how the token is matched. On the right side, the action to take. The value of the token is stored in `yylval` and the type of token is returned.

More details in the [Flex manual](#).

A Ruby equivalent, using the `rexical` gem (a port of Lex to Ruby), would be:

```
1 macro
2   BLANK      [\ \t]+
3
4 rule
5   # Whitespace
6   {BLANK}    # ignore
7
8   # Literals
9   [0-9]+     { [:NUMBER, text.to_i] }
```

```
10
11 # Keywords
12 end           { [:END, text] }
```

Rexical follows a similar grammar as Lex. Regular expression on the left and action on the right. However, an array of two items is used to return the type and value of the matched token.

Mode details on the [rexical project page](#).

RAGEL

A powerful tool for creating a scanner is [Ragel](#). It's described as a *State Machine Compiler*: lexers, like regular expressions, are state machines. Being very flexible, they can handle grammars of varying complexities and output parser in several languages.

Here's what a Ragel grammar looks like:

```
1  %%{
2  machine lexer;
3
4  # Machine
5  number      = [0-9]+;
6  whitespace  = " ";
7  keyword     = "end" | "def" | "class" | "if" | "else" | "true" | "false" | "nil";
8
9  # Actions
10 main := |*
11   whitespace; # ignore
12   number      => { tokens << [:NUMBER, data[ts..te].to_i] };
13   keyword     => { tokens << [data[ts...te].upcase.to_sym, data[ts...te]] };
14 *|;
15
16 class Lexer
17   def initialize
```

```
18     %% write data;
19 end
20
21 def run(data)
22     eof = data.size
23     line = 1
24     tokens = []
25     %% write init;
26     %% write exec;
27     tokens
28 end
29 end
30 }%%
```

More details in the [Ragel manual \(PDF\)](#).

Here are a few real-world examples of Ragel grammars used as language lexers:

- [Min's lexer](#) (Java)
- [Potion's lexer](#) (C)

PYTHON STYLE INDENTATION FOR AWESOME

If you intend to build a fully-functioning language, you should use one of the two previous tools. Since Awesome is a simplistic language and we want to illustrate the basic concepts of a scanner, we will build the lexer from scratch using regular expressions.

To make things more interesting, we'll use indentation to delimit blocks in our toy language, as in Python. All of indentation magic takes place within the lexer. Parsing blocks of code delimited with { . . . } is no different from parsing indentation when you know how to do it.

Tokenizing the following Python code:

```
1 if tasty == True:
2     print "Delicious!"
```

will yield these tokens:

```
1 [IDENTIFIER if] [IDENTIFIER tasty] [EQUAL] [IDENTIFIER True]
2 [INDENT] [IDENTIFIER print] [STRING "Delicious!"]
3 [DEDENT]
```

The block is wrapped in `INDENT` and `DEDENT` tokens instead of `{` and `}`.

The indentation-parsing algorithm is simple. You need to track two things: the current indentation level and the stack of indentation levels. When you encounter a line break followed by spaces, you update the indentation level. Here's our lexer for the Awesome language:

```
1 class Lexer
2     KEYWORDS = ["def", "class", "if", "true", "false", "nil"]
3
4     def tokenize(code)
5         # Cleanup code by remove extra line breaks
6         code.chomp!
7
8         # Current character position we're parsing
9         i = 0
10
11        # Collection of all parsed tokens in the form [:TOKEN_TYPE, value]
12        tokens = []
13
14        # Current indent level is the number of spaces in the last indent.
15        current_indent = 0
16        # We keep track of the indentation levels we are in so that when we dedent, we can
17        # check if we're on the correct level.
18        indent_stack = []
19
20        # This is how to implement a very simple scanner.
21        # Scan one character at the time until you find something to parse.
```

lexer.rb

```

22 while i < code.size
23   chunk = code[i..-1]
24
25   # Matching standard tokens.
26   #
27   # Matching if, print, method names, etc.
28   if identifier = chunk[/\A([a-z]\w*)/, 1]
29     # Keywords are special identifiers tagged with their own name, 'if' will result
30     # in an [:IF, "if"] token
31     if KEYWORDS.include?(identifier)
32       tokens << [identifier.upcase.to_sym, identifier]
33     # Non-keyword identifiers include method and variable names.
34     else
35       tokens << [:IDENTIFIER, identifier]
36     end
37     # skip what we just parsed
38     i += identifier.size
39
40     # Matching class names and constants starting with a capital letter.
41     elsif constant = chunk[/\A([A-Z]\w*)/, 1]
42       tokens << [:CONSTANT, constant]
43       i += constant.size
44
45     elsif number = chunk[/\A([0-9]+)/, 1]
46       tokens << [:NUMBER, number.to_i]
47       i += number.size
48
49     elsif string = chunk[/\A"(.*)" /, 1]
50       tokens << [:STRING, string]
51       i += string.size + 2
52
53     # Here's the indentation magic!
54     #
55     # We have to take care of 3 cases:
56     #
57     #   if true: # 1) the block is created
58     #     line 1
59     #     line 2 # 2) new line inside a block
60     #   continue # 3) dedent
61     #
62     # This elsif takes care of the first case. The number of spaces will determine
63     # the indent level.

```



```

64 elsif indent = chunk[/\A:\n( +)/m, 1] # Matches ": <newline> <spaces>"
65     # When we create a new block we expect the indent level to go up.
66     if indent.size <= current_indent
67         raise "Bad indent level, got #{indent.size} indents, " +
68             "expected > #{current_indent}"
69     end
70     # Adjust the current indentation level.
71     current_indent = indent.size
72     indent_stack.push(current_indent)
73     tokens << [:INDENT, indent.size]
74     i += indent.size + 2
75
76 # This elsif takes care of the two last cases:
77 # Case 2: We stay in the same block if the indent level (number of spaces) is the
78 #         same as current_indent.
79 # Case 3: Close the current block, if indent level is lower than current_indent.
80 elsif indent = chunk[/\A\n( *)/m, 1] # Matches "<newline> <spaces>"
81     if indent.size == current_indent # Case 2
82         # Nothing to do, we're still in the same block
83         tokens << [:NEWLINE, "\n"]
84     elsif indent.size < current_indent # Case 3
85         while indent.size < current_indent
86             indent_stack.pop
87             current_indent = indent_stack.first || 0
88             tokens << [:DEDENT, indent.size]
89         end
90         tokens << [:NEWLINE, "\n"]
91     else # indent.size > current_indent, error!
92         # Cannot increase indent level without using ":", so this is an error.
93         raise "Missing ':'"
94     end
95     i += indent.size + 1
96
97 # Match long operators such as ||, &&, ==, !=, <= and >=.
98 # One character long operators are matched by the catch all `else` at the bottom.
99 elsif operator = chunk[/\A(\|\|&&|=|!=|<=|>=)/, 1]
100     tokens << [operator, operator]
101     i += operator.size
102
103 # Ignore whitespace
104 elsif chunk.match(/\A /)
105     i += 1

```

```

106
107     # Catch all single characters
108     # We treat all other single characters as a token. Eg.: ( ) , . ! + - <
109     else
110         value = chunk[0,1]
111         tokens << [value, value]
112         i += 1
113
114     end
115
116 end
117
118 # Close all open blocks
119 while indent = indent_stack.pop
120     tokens << [:DEDENT, indent_stack.first || 0]
121 end
122
123 tokens
124 end
125 end

```

You can test the lexer yourself by running the test file included with the book. Run `ruby -Itest test/lexer_test.rb` from the code directory and it should output 0 failures, 0 errors. Here's an excerpt from that test file.

```

1 code = <<-CODE
2 if 1:
3   print "..."
4   if false:
5     pass
6   print "done!"
7 print "The End"
8 CODE
9 tokens = [
10  [:IF, "if"], [:NUMBER, 1],
11  [:INDENT, 2],
12  [:IDENTIFIER, "print"], [:STRING, "..."], [:NEWLINE, "\n"],
13  [:IF, "if"], [:FALSE, "false"],
14  [:INDENT, 4],

```

test/lexer_test.rb

```
15     [:IDENTIFIER, "pass"],
16     [:DEDENT, 2], [:NEWLINE, "\n"],
17     [:IDENTIFIER, "print"],
18     [:STRING, "done!"],
19     [:DEDENT, 0], [:NEWLINE, "\n"],
20     [:IDENTIFIER, "print"], [:STRING, "The End"]
21 ]
22 assert_equal tokens, Lexer.new.tokenize(code)
```

Some parsers take care of both lexing and parsing in their grammar. We'll see more about those in the next section.

DO IT YOURSELF I

- a. Modify the lexer to parse: `while condition: ...` control structures.
- b. Modify the lexer to delimit blocks with `{ ... }` instead of indentation.

[Solutions to Do It Yourself I.](#)

PARSER

By themselves, the tokens output by the lexer are just building blocks. The parser contextualizes them by organizing them in a structure. The lexer produces an array of tokens; the parser produces a tree of nodes.

Lets take those tokens from previous section:


```
1 [IDENTIFIER print] [STRING "I ate"] [COMMA]
2     [NUMBER 3] [COMMA]
3     [IDENTIFIER pies]
```

The most common parser output is an Abstract Syntax Tree, or AST. It's a tree of nodes that represents what the code means to the language. The previous lexer tokens will produce the following:

```
1 [<Call name=print,
2     arguments=[<String value="I ate">,
3         <Number value=3>,
4         <Local name=pies>]
5 >]
```

Or as a visual tree:

Figure 2



The parser found that `print` was a method call and the following tokens are the arguments.

Parser generators are commonly used to accomplish the otherwise tedious task of building a parser. Much like the English language, a programming language needs a grammar to define its rules. The parser generator will convert this grammar into a parser that will compile lexer tokens into AST nodes.

BISON (YACC)

Bison is a modern version of Yacc, the most widely used parser. Yacc stands for Yet Another Compiler Compiler, because it compiles the grammar to a compiler of tokens. It's used in several mainstream languages, like Ruby. Most often used with Lex, it has been ported to several target languages.

- [Racc for Ruby](#)
- [Ply for Python](#)
- [JavaCC for Java](#)

Like Lex, from the previous chapter, Yacc compiles a grammar into a parser. Here's how a Yacc grammar rule is defined:

```
1 Call: /* Name of the rule */
2   Expression '.' IDENTIFIER           { $$ = CallNode_new($1, $3, NULL); }
3 | Expression '.' IDENTIFIER '(' ArgList ')' { $$ = CallNode_new($1, $3, $5); }
4 /*  $1      $2      $3      $4      $5  $6  <= values from the rule are stored in
5                                     these variables. */
6 ;
```

On the left is defined how the rule can be matched using tokens and other rules.

On the right side, between brackets is the action to execute when the rule matches.

In that block, we can reference tokens being matched using \$1, \$2, etc. Finally, we store the result in \$\$.

LEMON

[Lemon](#) is quite similar to Yacc, with a few differences. From its website:

- Using a different grammar syntax which is less prone to programming errors.
- The parser generated by Lemon is both re-entrant and thread-safe.
- Lemon includes the concept of a non-terminal destructor, which makes it much easier to write a parser that does not leak memory.

For more information, refer to the [the manual](#) or check real examples inside [Potion](#).

ANTLR

[ANTLR](#) is another parsing tool. This one let's you declare lexing and parsing rules in the same grammar. It has been ported to [several target languages](#).

PEGS

Parsing Expression Grammars, or PEGs, are very powerful at parsing complex languages. I've used a PEG generated from [peg/leg](#) in tinyrb to parse Ruby's infamous syntax with encouraging results ([tinyrb's grammar](#)).

[Treetop](#) is an interesting Ruby tool for creating PEG.

OPERATOR PRECEDENCE

One of the common pitfalls of language parsing is operator precedence. Parsing $x + y * z$ should not produce the same result as $(x + y) * z$, same for all other

operators. Each language has an operator precedence table, often based on mathematics order of operations. Several ways to handle this exist. Yacc-based parsers implement the [Shunting Yard algorithm](#) in which you give a precedence level to each kind of operator. Operators are declared in Bison and Yacc with `%left` and `%right` macros. Read more in [Bison's manual](#).

Here's the operator precedence table for our language, based on the [C language operator precedence](#):

```
1 left  '.'
2 right '!'
3 left  '*' '/'
4 left  '+' '-'
5 left  '>' '>=' '<' '<='
6 left  '==' '!='
7 left  '&&'
8 left  '||'
9 right '='
10 left  ','
```

The higher the precedence (top is higher), the sooner the operator will be parsed. If the line `a + b * c` is being parsed, the part `b * c` will be parsed first since `*` has higher precedence than `+`. Now, if several operators having the same precedence are competing to be parsed all the once, the conflict is resolved using associativity, declared with the `left` and `right` keyword before the token. For example, with the expression `a = b = c`. Since `=` has right-to-left associativity, it will start parsing from the right, `b = c`. Resulting in `a = (b = c)`.

For other types of parsers (ANTLR and PEG) a simpler but less efficient alternative can be used. Simply declaring the grammar rules in the right order will produce the desired result:

```
1 expression:      equality
2 equality:        additive ( ( '=' | '!=' ) additive )*
3 additive:       multiplicative ( ( '+' | '-' ) multiplicative )*
4 multiplicative: primary ( ( '*' | '/' ) primary )*
5 primary:        '(' expression ')' | NUMBER | VARIABLE | '-' primary
```

The parser will try to match rules recursively, starting from `expression` and finding its way to `primary`. Since `multiplicative` is the last rule called in the parsing process, it will have greater precedence.

CONNECTING THE LEXER AND PARSER IN AWESOME

For our Awesome parser we'll use Racc, the Ruby version of Yacc. It's much harder to build a parser from scratch than it is to create a lexer. However, most languages end up writing their own parser because the result is faster and provides better error reporting.

The input file you supply to Racc contains the grammar of your language and is very similar to a Yacc grammar.

```
1 class Parser
2
3 # Declare tokens produced by the lexer
4 token IF ELSE
5 token DEF
6 token CLASS
7 token NEWLINE
8 token NUMBER
9 token STRING
10 token TRUE FALSE NIL
11 token IDENTIFIER
12 token CONSTANT
13 token INDENT DEDENT
14
```

grammar.y

```

15 # Precedence table
16 # Based on http://en.wikipedia.org/wiki/Operators_in_C_and_C%2B%2B#Operator_precedence
17 prechigh
18   left  '.'
19   right '!'
20   left  '*' '/'
21   left  '+' '-'
22   left  '>' '>=' '<' '<='
23   left  '==' '!='
24   left  '&&'
25   left  '||'
26   right '='
27   left  ','
28 preclow
29
30 rule
31   # All rules are declared in this format:
32   #
33   #   RuleName:
34   #       OtherRule TOKEN AnotherRule    { code to run when this matches }
35   #   | OtherRule                        { ... }
36   #   ;
37   #
38   # In the code section (inside the {...} on the right):
39   # - Assign to "result" the value returned by the rule.
40   # - Use val[index of expression] to reference expressions on the left.
41
42
43 # All parsing will end in this rule, being the trunk of the AST.
44 Root:
45   /* nothing */                { result = Nodes.new([]) }
46   | Expressions                 { result = val[0] }
47   ;
48
49 # Any list of expressions, class or method body, seperated by line breaks.
50 Expressions:
51   Expression                    { result = Nodes.new(val) }
52   | Expressions Terminator Expression { result = val[0] << val[2] }
53   # To ignore trailing line breaks
54   | Expressions Terminator      { result = val[0] }
55   | Terminator                  { result = Nodes.new([]) }
56   ;

```

```

57
58 # All types of expressions in our language
59 Expression:
60     Literal
61     | Call
62     | Operator
63     | Constant
64     | Assign
65     | Def
66     | Class
67     | If
68     | '(' Expression ')'      { result = val[1] }
69 ;
70
71 # All tokens that can terminate an expression
72 Terminator:
73     NEWLINE
74     | ";"
75 ;
76
77 # All hard-coded values
78 Literal:
79     NUMBER                { result = NumberNode.new(val[0]) }
80     | STRING              { result = StringNode.new(val[0]) }
81     | TRUE                { result = TrueNode.new }
82     | FALSE              { result = FalseNode.new }
83     | NIL                { result = NilNode.new }
84 ;
85
86 # A method call
87 Call:
88     # method
89     IDENTIFIER            { result = CallNode.new(nil, val[0], []) }
90     # method(arguments)
91     | IDENTIFIER "(" ArgList ")" { result = CallNode.new(nil, val[0], val[2]) }
92     # receiver.method
93     | Expression "." IDENTIFIER { result = CallNode.new(val[0], val[2], []) }
94     # receiver.method(arguments)
95     | Expression "."
96     IDENTIFIER "(" ArgList ")" { result = CallNode.new(val[0], val[2], val[4]) }
97 ;
98

```

```

99 ArgList:
100     /* nothing */           { result = [] }
101 | Expression               { result = val }
102 | ArgList "," Expression   { result = val[0] << val[2] }
103 ;
104
105 Operator:
106 # Binary operators
107 Expression '||' Expression { result = CallNode.new(val[0], val[1], [val[2]]) }
108 | Expression '&&' Expression { result = CallNode.new(val[0], val[1], [val[2]]) }
109 | Expression '==' Expression { result = CallNode.new(val[0], val[1], [val[2]]) }
110 | Expression '!=' Expression { result = CallNode.new(val[0], val[1], [val[2]]) }
111 | Expression '>' Expression { result = CallNode.new(val[0], val[1], [val[2]]) }
112 | Expression '>=' Expression { result = CallNode.new(val[0], val[1], [val[2]]) }
113 | Expression '<' Expression { result = CallNode.new(val[0], val[1], [val[2]]) }
114 | Expression '<=' Expression { result = CallNode.new(val[0], val[1], [val[2]]) }
115 | Expression '+' Expression { result = CallNode.new(val[0], val[1], [val[2]]) }
116 | Expression '-' Expression { result = CallNode.new(val[0], val[1], [val[2]]) }
117 | Expression '*' Expression { result = CallNode.new(val[0], val[1], [val[2]]) }
118 | Expression '/' Expression { result = CallNode.new(val[0], val[1], [val[2]]) }
119 ;
120
121 Constant:
122     CONSTANT                 { result = GetConstantNode.new(val[0]) }
123 ;
124
125 # Assignment to a variable or constant
126 Assign:
127     IDENTIFIER "=" Expression { result = SetLocalNode.new(val[0], val[2]) }
128 | CONSTANT "=" Expression    { result = SetConstantNode.new(val[0], val[2]) }
129 ;
130
131 # Method definition
132 Def:
133     DEF IDENTIFIER Block      { result = DefNode.new(val[1], [], val[2]) }
134 | DEF IDENTIFIER
135     "(" ParamList ")" Block  { result = DefNode.new(val[1], val[3], val[5]) }
136 ;
137
138 ParamList:
139     /* nothing */           { result = [] }
140 | IDENTIFIER                 { result = val }

```

```

141 | ParamList "," IDENTIFIER      { result = val[0] << val[2] }
142 ;
143
144 # Class definition
145 Class:
146   CLASS CONSTANT Block        { result = ClassNode.new(val[1], val[2]) }
147 ;
148
149 # if block
150 If:
151   IF Expression Block          { result = IfNode.new(val[1], val[2]) }
152 ;
153
154 # A block of indented code. You see here that all the hard work was done by the
155 # lexer.
156 Block:
157   INDENT Expressions DEDENT     { result = val[1] }
158 # If you don't like indentation you could replace the previous rule with the
159 # following one to separate blocks w/ curly brackets. You'll also need to remove the
160 # indentation magic section in the lexer.
161 # "{" Expressions "}"          { replace = val[1] }
162 ;
163 end
164
165 ---- header
166   require "lexer"
167   require "nodes"
168
169 ---- inner
170 # This code will be put as-is in the Parser class.
171 def parse(code, show_tokens=false)
172   @tokens = Lexer.new.tokenize(code) # Tokenize the code using our lexer
173   puts @tokens.inspect if show_tokens
174   do_parse # Kickoff the parsing process
175 end
176
177 def next_token
178   @tokens.shift
179 end

```


We then generate the parser with: `racc -o parser.rb grammar.y`. This will create a `Parser` class that we can use to parse our code. Run `ruby -Itest test/parser_test.rb` from the code directory to test the parser. Here's an excerpt from this file.

```
1 code = <<-CODE
2 def method(a, b):
3   true
4 CODE
5
6 nodes = Nodes.new([
7   DefNode.new("method", ["a", "b"]),
8   Nodes.new([TrueNode.new])
9 ])
10 ])
11
12 assert_equal nodes, Parser.new.parse(code)
```

test/parser_test.rb

Parsing code will return a tree of nodes. The root node will always be of type `Nodes` which contains children nodes.

DO IT YOURSELF II

- a. Add a rule in the grammar to parse `while` blocks.
- b. Add a grammar rule to handle the `!` unary operators, eg.: `!x`. Making the following test pass (`test_unary_operator`):

[Solutions to Do It Yourself II.](#)

RUNTIME MODEL

The runtime model of a language is how we represent its objects, its methods, its types, its structure in memory. If the parser determines how you *talk* to the language, the runtime defines how the language *behaves*. Two languages could share the same parser but have different runtimes and be very different.

When designing your runtime, there are three factors you will want to consider:

- Speed: most of the speed will be due to the efficiency of the runtime.
- Flexibility: the more you allow the user to modify the language, the more powerful it is.
- Memory footprint: of course, all of this while using as little memory as possible.

As you might have noticed, these three constraints are mutually conflictual. Designing a language is always a game of give-and-take.

With those considerations in mind, there are several ways you can model your runtime.

PROCEDURAL

One of the simplest runtime models, like C and PHP (before version 4). Everything is centered around methods (procedures). There aren't any objects and all methods often share the same namespace. It gets messy pretty quickly!

CLASS-BASED

The class-based model is the most popular at the moment. Think of Java, Python, Ruby, etc. It might be the easiest model to understand for the users of your language.

PROTOTYPE-BASED

Except for Javascript, no Prototype-based languages have reached widespread popularity yet. This model is the easiest one to implement and also the most flexible because everything is a clone of an object.

Ian Piumarta describes how to design an [Open, Extensible Object Model](#) that allows the language's users to modify its behavior at runtime.

Look at the appendix at the end of this book for a sample prototype-based language: [Appendix: Mio, a minimalist homoiconic language](#).

FUNCTIONAL

The functional model, used by Lisp and other languages, treats computation as the evaluation of mathematical functions and avoids state and mutable data. This model has its roots in Lambda Calculus.

OUR AWESOME RUNTIME

Since most of us are familiar with Class-based runtimes, I decided to use that for our Awesome language. The following code defines how objects, methods and classes are stored and how they interact together.

The `AwesomeObject` class is the central object of our runtime. Since everything is an object in our language, everything we will put in the runtime needs to be an

object, thus an instance of this class. `AwesomeObject`s have a class and can hold a ruby value. This will allow us to store data such as a string or a number in an object to keep track of its Ruby representation.

```
1 # Represents an Awesome object instance in the Ruby world.
2 class AwesomeObject
3   attr_accessor :runtime_class, :ruby_value
4
5   # Each object have a class (named runtime_class to prevent errors with Ruby's class
6   # method). Optionally an object can hold a Ruby value (eg.: numbers and strings).
7   def initialize(runtime_class, ruby_value=self)
8     @runtime_class = runtime_class
9     @ruby_value = ruby_value
10  end
11
12  # Call a method on the object.
13  def call(method, arguments=[])
14    # Like a typical Class-based runtime model, we store methods in the class of the
15    # object.
16    @runtime_class.lookup(method).call(self, arguments)
17  end
18 end
```

runtime/object.rb

Remember that in `Awesome`, everything is an object. Even classes are instances of the `Class` class. `AwesomeClasses` hold the methods and can be instantiated via their `new` method.

```
1 # Represents a Awesome class in the Ruby world. Classes are objects in Awesome so they
2 # inherit from AwesomeObject.
3 class AwesomeClass < AwesomeObject
4   attr_reader :runtime_methods
5
6   # Creates a new class. Number is an instance of Class for example.
7   def initialize
8     @runtime_methods = {}
9
10    # Check if we're bootstrapping (launching the runtime). During this process the
```

runtime/class.rb

```

11 # runtime is not fully initialized and core classes do not yet exists, so we defer
12 # using those once the language is bootstrapped.
13 # This solves the chicken-or-the-egg problem with the Class class. We can
14 # initialize Class then set Class.class = Class.
15 if defined?(Runtime)
16   runtime_class = Runtime["Class"]
17 else
18   runtime_class = nil
19 end
20
21 super(runtime_class)
22 end
23
24 # Lookup a method
25 def lookup(method_name)
26   method = @runtime_methods[method_name]
27   unless method
28     raise "Method not found: #{method_name}"
29   end
30   method
31 end
32
33 # Create a new instance of this class
34 def new
35   AwesomeObject.new(self)
36 end
37
38 # Create an instance of this Awesome class that holds a Ruby value. Like a String,
39 # Number or true.
40 def new_with_value(value)
41   AwesomeObject.new(self, value)
42 end
43 end

```

And here's the method object which will store methods defined from within our runtime.

```

1 # Represents a method defined in the runtime.
2 class AwesomeMethod
3   def initialize(params, body)

```

runtime/method.rb

```

4     @params = params
5     @body = body
6 end
7
8 def call(receiver, arguments)
9     # Create a context of evaluation in which the method will execute.
10    context = Context.new(receiver)
11
12    # Assign arguments to local variables
13    @params.each_with_index do |param, index|
14        context.locals[param] = arguments[index]
15    end
16
17    @body.eval(context)
18 end
19 end

```

Notice that we use the `call` method for evaluating a method. That will allow us to define runtime methods from Ruby using `Procs`. Here's why:

```

1 p = proc do |arg1, arg2|
2     # ...
3 end
4 p.call(1, 2) # execute the block of code passed to proc (everything between do ... end)

```

`Procs` can be executed via their `call` method. We'll see how this is used to defined runtime methods from Ruby in the bootstrapping section of the runtime.

Before we bootstrap our runtime, there is one missing object we need to define and that is the context of evaluation. The `Context` object encapsulates the environment of evaluation of a specific block of code. It will keep track of the following:

- Local variables.
- The current value of `self`, the object on which methods with no receivers are called, eg.: `print` is like `self.print`.

- The current class, the class on which methods are defined with the `def` keyword.

This is also where our constants (ie. classes) will be stored.

```
1 # The evaluation context.
2 class Context
3   attr_reader :locals, :current_self, :current_class
4
5   # We store constants as class variable (class variables start with @@ and instance
6   # variables start with @ in Ruby) since they are globally accessible. If you want to
7   # implement namespacing of constants, you could store it in the instance of this
8   # class.
9   @@constants = {}
10
11  def initialize(current_self, current_class=current_self.runtime_class)
12    @locals = {}
13    @current_self = current_self
14    @current_class = current_class
15  end
16
17  # Shortcuts to access constants, Runtime[...] instead of Runtime.constants[...]
18  def [](name)
19    @@constants[name]
20  end
21  def []=(name, value)
22    @@constants[name] = value
23  end
24 end
```

runtime/context.rb

Finally, we bootstrap the runtime. At first, no objects exist in the runtime. Before we can execute our first expression, we need to populate that runtime with a few objects: `Class`, `Object`, `true`, `false`, `nil` and a few core methods.

```
1 # Bootstrap the runtime. This is where we assemble all the classes and objects together
2 # to form the runtime.
3
4                                     # What's happening in the runtime:
```

runtime/bootstrap.rb

```

4 awesome_class = AwesomeClass.new          # Class
5 awesome_class.runtime_class = awesome_class # Class.class = Class
6 object_class = AwesomeClass.new          # Object = Class.new
7 object_class.runtime_class = awesome_class # Object.class = Class
8
9 # Create the Runtime object (the root context) on which all code will start its
10 # evaluation.
11 Runtime = Context.new(object_class.new)
12
13 Runtime["Class"] = awesome_class
14 Runtime["Object"] = object_class
15 Runtime["Number"] = AwesomeClass.new
16 Runtime["String"] = AwesomeClass.new
17
18 # Everything is an object in our language, even true, false and nil. So they need
19 # to have a class too.
20 Runtime["TrueClass"] = AwesomeClass.new
21 Runtime["FalseClass"] = AwesomeClass.new
22 Runtime["NilClass"] = AwesomeClass.new
23
24 Runtime["true"] = Runtime["TrueClass"].new_with_value(true)
25 Runtime["false"] = Runtime["FalseClass"].new_with_value(false)
26 Runtime["nil"] = Runtime["NilClass"].new_with_value(nil)
27
28 # Add a few core methods to the runtime.
29
30 # Add the `new` method to classes, used to instantiate a class:
31 # eg.: Object.new, Number.new, String.new, etc.
32 Runtime["Class"].runtime_methods["new"] = proc do |receiver, arguments|
33   receiver.new
34 end
35
36 # Print an object to the console.
37 # eg.: print("hi there!")
38 Runtime["Object"].runtime_methods["print"] = proc do |receiver, arguments|
39   puts arguments.first.ruby_value
40   Runtime["nil"]
41 end

```

Now that we got all the pieces together we can call methods and create objects inside our runtime.


```
1 # Mimic Object.new in the language
2 object = Runtime["Object"].call("new")
3
4 assert_equal Runtime["Object"], object.runtime_class # assert object is an Object
```

Can you feel the language coming alive? We'll learn how to map that runtime to the nodes we created from our parser in the next section.

DO IT YOURSELF III

- a. Implement inheritance by adding a superclass to each Awesome class.
- b. Add the method to handle $x + 2$.

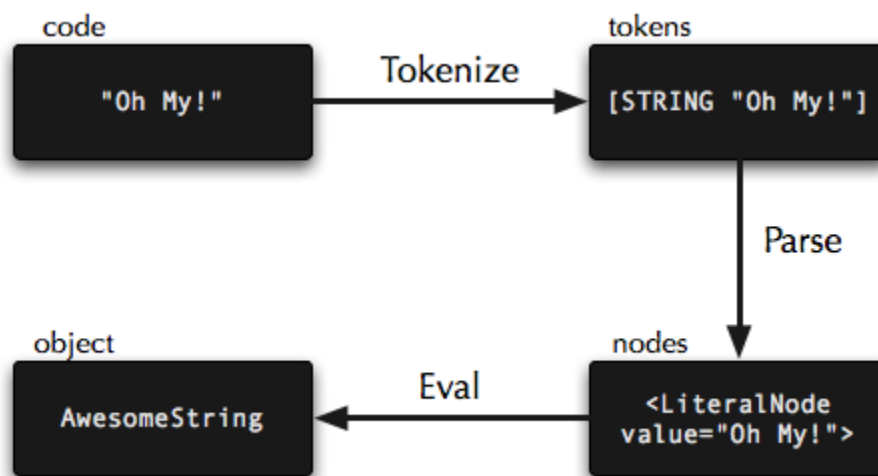
[Solutions to Do It Yourself III.](#)

INTERPRETER

The interpreter is the module that evaluates the code. It reads the AST produced by the parser and executes each action associated with the nodes, modifying the runtime.

Figure 3 recapitulates the path of a string in our language.

Figure 3



The lexer creates the token, the parser takes those tokens and converts it into nodes. Finally, the interpreter evaluates the nodes.

A common approach to execute an AST is to implement a [Visitor class](#) that visits all the nodes one by one, running the appropriate code. This make things even more modular and eases the optimization efforts on the AST. But for the purpose of this book, we'll keep things simple and let each node handle its evaluation.

Remember the nodes we created in the parser: `StringNode` for a string, `ClassNode` for a class definition? Here we're reopening those classes and adding a new method to each one: `eval`. This method will be responsible for interpreting that particular node.

```
1 require "parser"
2 require "runtime"
3
4 class Interpreter
5   def initialize
6     @parser = Parser.new
7   end
8
9   def eval(code)
10    @parser.parse(code).eval(Runtime)
11  end
12 end
13
14 class Nodes
15   # This method is the "interpreter" part of our language. All nodes know how to eval
16   # itself and returns the result of its evaluation by implementing the "eval" method.
17   # The "context" variable is the environment in which the node is evaluated (local
18   # variables, current class, etc.).
19   def eval(context)
20     return_value = nil
21     nodes.each do |node|
22       return_value = node.eval(context)
23     end
24     # The last value evaluated in a method is the return value. Or nil if none.
25     return_value || Runtime["nil"]
26   end
27 end
28
29 class NumberNode
30   def eval(context)
31     # Here we access the Runtime, which we'll see in the next section, to create a new
32     # instance of the Number class.
33     Runtime["Number"].new_with_value(value)
34   end
35 end
36
37 class StringNode
38   def eval(context)
39     Runtime["String"].new_with_value(value)
40   end
41 end
42
```

```

43 class TrueNode
44   def eval(context)
45     Runtime["true"]
46   end
47 end
48
49 class FalseNode
50   def eval(context)
51     Runtime["false"]
52   end
53 end
54
55 class NilNode
56   def eval(context)
57     Runtime["nil"]
58   end
59 end
60
61 class CallNode
62   def eval(context)
63     # If there's no receiver and the method name is the name of a local variable, then
64     # it's a local variable access. This trick allows us to skip the () when calling a
65     # method.
66     if receiver.nil? && context.locals[method] && arguments.empty?
67       context.locals[method]
68
69     # Method call
70     else
71       if receiver
72         value = receiver.eval(context)
73       else
74         # In case there's no receiver we default to self, calling "print" is like
75         # "self.print".
76         value = context.current_self
77       end
78
79       eval_arguments = arguments.map { |arg| arg.eval(context) }
80       value.call(method, eval_arguments)
81     end
82   end
83 end
84

```

```

85 class GetConstantNode
86   def eval(context)
87     context[name]
88   end
89 end
90
91 class SetConstantNode
92   def eval(context)
93     context[name] = value.eval(context)
94   end
95 end
96
97 class SetLocalNode
98   def eval(context)
99     context.locals[name] = value.eval(context)
100  end
101 end
102
103 class DefNode
104   def eval(context)
105     # Defining a method is adding a method to the current class.
106     method = AwesomeMethod.new(params, body)
107     context.current_class.runtime_methods[name] = method
108   end
109 end
110
111 class ClassNode
112   def eval(context)
113     # Try to locate the class. Allows reopening classes to add methods.
114     awesome_class = context[name]
115
116     unless awesome_class # Class doesn't exist yet
117       awesome_class = AwesomeClass.new
118       # Register the class as a constant in the runtime.
119       context[name] = awesome_class
120     end
121
122     # Evaluate the body of the class in its context. Providing a custom context allows
123     # to control where methods are added when defined with the def keyword. In this
124     # case, we add them to the newly created class.
125     class_context = Context.new(awesome_class, awesome_class)
126

```

```

127     body.eval(class_context)
128
129     awesome_class
130 end
131 end
132
133 class IfNode
134   def eval(context)
135     # We turn the condition node into a Ruby value to use Ruby's "if" control
136     # structure.
137     if condition.eval(context).ruby_value
138       body.eval(context)
139     end
140   end
141 end

```

The interpreter part (the `eval` method) is the connector between the parser and the runtime of our language. Once we call `eval` on the root node, all children nodes are evaluated recursively. This is why we call the output of the parser an AST, for Abstract Syntax Tree. It is a tree of nodes. And evaluating the top level node of that tree will have the cascading effect of evaluating each of its children.

Let's run our first full program!

```

1 code = <<-CODE
2 class Awesome:
3   def does_it_work:
4     "yeah!"
5
6 awesome_object = Awesome.new
7 if awesome_object:
8   print(awesome_object.does_it_work)
9 CODE
10
11 assert_prints("yeah!\n") { Interpreter.new.eval(code) }

```

test/interpreter_test.rb

To complete our language we can create a script to run a file or an REPL (for read-eval-print-loop), or interactive interpreter.

awesome

```
1 #!/usr/bin/env ruby
2 # The Awesome language!
3 #
4 # usage:
5 #   ./awesome example.awm # to eval a file
6 #   ./awesome             # to start the REPL
7 #
8 # on Windows run with: ruby awesome [options]
9
10 $:.unshift "." # Fix for Ruby 1.9
11 require "interpreter"
12 require "readline"
13
14 interpreter = Interpreter.new
15
16 # If a file is given we eval it.
17 if file = ARGV.first
18   interpreter.eval File.read(file)
19
20 # Start the REPL, read-eval-print-loop, or interactive interpreter
21 else
22   puts "Awesome REPL, CTRL+C to quit"
23   loop do
24     line = Readline::readline(">> ")
25     Readline::HISTORY.push(line)
26     value = interpreter.eval(line)
27     puts "=> #{value.ruby_value.inspect}"
28   end
29
30 end
```

Run the interactive interpreter by running `./awesome` and type a line of Awesome code, eg.: `print("It works!")`. Here's a sample Awesome session:

```
1 Awesome REPL, CTRL+C to quit
2 >> m = "This is Awesome!"
```

```
3 => "This is Awesome!"
4 >> print(m)
5 This is Awesome!
6 => nil
```

Also, try running a file: `./awesome example.awm`.

DO IT YOURSELF IV

- a. Implement the `WhileNode eval` method.

[Solutions to Do It Yourself IV.](#)

COMPILATION

Dynamic languages like Ruby and Python have many advantages, like development speed and flexibility. But sometimes, compiling a language to a more efficient format might be more appropriate. When you write C code, you compile it to machine code. Java is compiled to a specific byte-code that you can run on the Java Virtual Machine (JVM).

These days, most dynamic languages also compile source code to machine code or byte-code on the fly, this is called Just In Time (JIT) compilation. It yields faster execution because executing code by walking an AST of nodes (like in Awesome) is less efficient than running machine code or byte-code. The reason why byte-code execution is faster is because it's closer to machine code than an AST is. Bringing your execution model closer to the machine will always produce faster results.

If you want to compile your language you have multiple choices, you can:

- compile to your own byte-code format and create a Virtual Machine to run it,
- compile it to JVM byte-code using tools like [ASM](#),
- write your own machine code compiler,
- use an existing compiler framework like [LLVM](#).

Compiling to a custom byte-code format is the most popular option for dynamic languages. Python and new Ruby implementations are going this way. But, this only makes sense if you're coding in a low level language like C. Because Virtual Machines are very simple (it's basically a loop and switch-case) you need total control over the structure of your program in memory to make it as fast as possible.

If you want more details on Virtual Machines, jump to the next chapter.

USING LLVM FROM RUBY

We'll be using LLVM Ruby bindings to compile a subset of Awesome to machine code on the fly. To compile a full language, most parts of the runtime have to be rewritten to be accessible from inside LLVM.

First, you'll need to install LLVM and the Ruby bindings. You can find instructions on [ruby-llvm project page](#). Installing LLVM can take quite some time, make sure you got yourself a tasty beverage before launching compilation.

Here's how to use LLVM from Ruby.

```
1 # Creates a new module to hold the code
2 mod = LLVM::Module.create("awesome")
3 # Creates the main function that will be called
4 main = mod.functions.add("main", [INT, LLVM::Type.pointer(PCHAR)], INT)
5 # Create a block of code to build machine code within
6 builder = LLVM::Builder.create
7 builder.position_at_end(main.basic_blocks.append)
8
9 # Find the function we're calling in the module
10 func = mod.functions.named("puts")
11 # Call the function
12 builder.call(func, builder.global_string_pointer("hello"))
13 # Return
14 builder.ret(LLVM::Int(0))
```

This is the equivalent of the following C code. In fact, it will generate similar machine code.

```
1 int main (int argc, char const *argv[]) {
2     puts("hello");
3     return 0;
4 }
```

The difference is that with LLVM we can dynamically generate the machine code. Because of that, we can create machine code from Awesome code.

COMPILING AWESOME TO MACHINE CODE

To compile Awesome to machine code, we'll create a `Compiler` class that will encapsulate the logic of calling LLVM to generate the byte-code. Then, we'll extend the nodes created by the parser to make them use the compiler.

```
1  require "rubygems"
2  require "parser"
3  require "nodes"
4
5  require 'llvm/core'
6  require 'llvm/execution_engine'
7  require 'llvm/transforms/scalar'
8  require 'llvm/transforms/ipo'
9
10 LLVM.init_x86
11
12 # Compiler is used in a similar way as the runtime. But, instead of executing code, it
13 # will generate LLVM byte-code for later execution.
14 class Compiler
15
16   # Initialize LLVM types
17   PCHAR = LLVM::Type.pointer(LLVM::Int8) # equivalent to *char in C
18   INT   = LLVM::Int # equivalent to int in C
19
20   attr_reader :locals
21
22   def initialize(mod=nil, function=nil)
23     # Create the LLVM module in which to store the code
24     @module = mod || LLVM::Module.create("awesome")
25
26     # To track local names during compilation
27     @locals = {}
28
29     # Function in which the code will be put
30     @function = function ||
```

compiler.rb

```

31         # By default we create a main function as it's the standard entry point
32         @module.functions.named("main") ||
33         @module.functions.add("main", [INT, LLVM::Type.pointer(PCHAR)], INT)
34
35     # Create an LLVM byte-code builder
36     @builder = LLVM::Builder.create
37     @builder.position_at_end(@function.basic_blocks.append)
38
39     @engine = LLVM::ExecutionEngine.create_jit_compiler(@module)
40 end
41
42 # Initial header to initialize the module.
43 def preamble
44     define_external_functions
45 end
46
47 def finish
48     @builder.ret(LLVM::Int(0))
49 end
50
51 # Create a new string.
52 def new_string(value)
53     @builder.global_string_pointer(value)
54 end
55
56 # Create a new number.
57 def new_number(value)
58     LLVM::Int(value)
59 end
60
61 # Call a function.
62 def call(func, args=[])
63     f = @module.functions.named(func)
64     @builder.call(f, *args)
65 end
66
67 # Assign a local variable
68 def assign(name, value)
69     # Allocate the memory and returns a pointer to it
70     ptr = @builder.alloca(value.type)
71     # Store the value insite the pointer
72     @builder.store(value, ptr)

```

```

73     # Keep track of the pointer so the compiler can find it back name later.
74     @locals[name] = ptr
75 end
76
77 # Load the value of a local variable.
78 def load(name)
79     @builder.load(@locals[name])
80 end
81
82 # Defines a function.
83 def function(name)
84     func = @module.functions.add(name, [], INT)
85     generator = Compiler.new(@module, func)
86     yield generator
87     generator.finish
88 end
89
90 # Optimize the generated LLVM byte-code.
91 def optimize
92     @module.verify!
93     pass_manager = LLVM::PassManager.new(@engine)
94     pass_manager.simplifycfg! # Simplify the CFG
95     pass_manager.mem2reg!     # Promote Memory to Register
96     pass_manager.gdce!       # Dead Global Elimination
97 end
98
99 # JIT compile and run the LLVM byte-code.
100 def run
101     @engine.run_function(@function, 0, 0)
102 end
103
104 def dump
105     @module.dump
106 end
107
108 private
109     def define_external_functions
110         fun = @module.functions.add("printf", [LLVM::Type.pointer(PCHAR)], INT, { :varargs => true })
111         fun.linkage = :external
112
113         fun = @module.functions.add("puts", [PCHAR], INT)
114         fun.linkage = :external

```

```

115
116     fun = @module.functions.add("read", [INT, PCHAR, INT], INT)
117     fun.linkage = :external
118
119     fun = @module.functions.add("exit", [INT], INT)
120     fun.linkage = :external
121     end
122 end
123
124 # Reopen class supported by the compiler to implement how each node is compiled
125 # (compile method).
126
127 class Nodes
128     def compile(compiler)
129         nodes.map { |node| node.compile(compiler) }.last
130     end
131 end
132
133 class NumberNode
134     def compile(compiler)
135         compiler.new_number(value)
136     end
137 end
138
139 class StringNode
140     def compile(compiler)
141         compiler.new_string(value)
142     end
143 end
144
145 class CallNode
146     def compile(compiler)
147         raise "Receiver not supported for compilation" if receiver
148
149         # Local variable access
150         if receiver.nil? && arguments.empty? && compiler.locals[method]
151             compiler.load(method)
152
153         # Method call
154         else
155             compiled_arguments = arguments.map { |arg| arg.compile(compiler) }
156             compiler.call(method, compiled_arguments)

```

```

157     end
158   end
159 end
160
161 class SetLocalNode
162   def compile(compiler)
163     compiler.assign(name, value.compile(compiler))
164   end
165 end
166
167 class DefNode
168   def compile(compiler)
169     raise "Parameters not supported for compilation" if !params.empty?
170     compiler.function(name) do |function|
171       body.compile(function)
172     end
173   end
174 end

```

With the compiler integrated with the nodes, we can now compile a simple program.

```

1 code = <<-CODE
2 def say_it:
3   x = "This is compiled!"
4   puts(x)
5 say_it
6 CODE
7
8 # Parse the code
9 node = Parser.new.parse(code)
10
11 # Compile it
12 compiler = Compiler.new
13 compiler.preamble
14 node.compile(compiler)
15 compiler.finish
16
17 # Uncomment to output LLVM byte-code
18 # compiler.dump

```

test/compiler_test.rb

```
19
20 # Optimize the LLVM byte-code
21 compiler.optimize
22
23 # JIT compile & execute
24 compiler.run
```

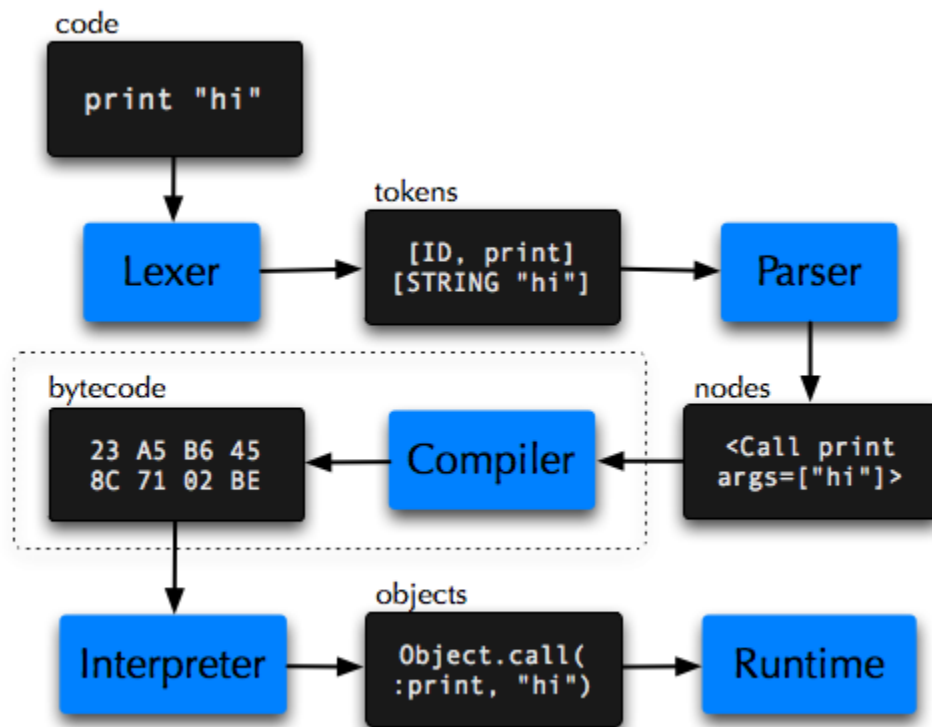
Remember that the compiler only supports a subset of our Awesome language. For example, object-oriented programming is not supported. To implement this, the runtime and structures used to store the classes and objects have to be loaded from inside the LLVM module. You can do this by compiling your runtime to LLVM byte-code, either writing it in C and using the C-to-LLVM compiler shipped with LLVM or by writing your runtime in a subset of your language that can be compiled to LLVM byte-code.

VIRTUAL MACHINE

If you're serious about speed and are ready to implement your language in C/C++, you need to introduce a VM (short for Virtual Machine) in your design.

When running your language on a VM you have to compile your AST nodes to byte-code. Despite adding an extra step, execution of the code is faster because byte-code format is closer to machine code than an AST. Figure 4 shows how you can introduce a byte-code compiler into your language design.

Figure 4




Look at `tinypy` and `tinyrb` in the Interesting Languages section for complete examples of small VM-based languages.

BYTE-CODE

The purpose of a VM is to bring the code as close to machine code as possible but without being too far from the original language making it hard (slow) to compile. An ideal byte-code is fast to compile from the source language and fast to execute while being very compact in memory.

The byte-code of a program consists of instructions. Each instruction starts with an opcode, specifying what this instruction does, followed by operands, which are the arguments of the instructions.

Figure 5



Most VMs share a similar set of instructions. Common ones include, `getlocal` to push a local variable value to the stack, `putstring` to push a string to the stack, `pop` to remove an item from the stack, `dup` to duplicate it, etc. You can see Ruby 1.9 (YARV) instruction set at lifegoo.pluskid.org. And the JVM instruction set at xs4all.nl.

TYPES OF VM

Stack-based virtual machines are most common and simpler to implement. Python, Ruby 1.9 and the JVM are all stack-based VMs. The execution of those is based around a stack in which values are stored to be passed around to instructions. Instructions will often pop values from the stack, modify it and push the result back to the stack.

Register-based VMs are becoming increasingly popular. Lua is register-based. They are closer to how a real machine work but produce bigger instructions with lots of operands but often fewer total instructions than a stack-based VM.

PROTOTYPING A VM IN RUBY

For the sake of understanding the inner working of a VM, we'll look at a prototype written in Ruby. It should be noted that, never in any case, a real VM should be written in a high level languages like Ruby. It defeats the purpose of bringing the code closer to the machine. Each line of code in a Virtual Machine is optimized to require as little machine instructions as possible to execute, high level language do not provide that control but C and C++ do.

```
1 # Bytecode
2 PUSH   = 0
3 ADD    = 1
4 PRINT  = 2
5 RETURN = 3
6
7 class VM
8   def run(bytecode)
9     # Stack to pass value between instructions.
10    stack = []
11    # Instruction Pointer, index of current instruction being executed in bytecode.
12    ip = 0
13
14    while true
15      case bytecode[ip]
16      when PUSH
17        stack.unshift bytecode[ip+=1]
18      when ADD
19        stack.unshift stack.pop + stack.pop
20      when PRINT
21        puts stack.pop
22      when RETURN
23        return
24    end
```

vm/vm.rb

```
25
26     # Continue to next instruction
27     ip += 1
28     end
29 end
30 end
31
32 VM.new.run [
33     # Here is the bytecode of our program, the equivalent of: print 1 + 2.
34     # Opcode, Operand      # Status of the stack after execution of the instruction.
35     PUSH,    1,           # stack = [1]
36     PUSH,    2,           # stack = [2, 1]
37     ADD,     ,            # stack = [3]
38     PRINT,   ,            # stack = []
39     RETURN
40 ]
```

As you can see, a Virtual Machine is simply a loop and a `switch-case`.

The byte-code that we execute is the equivalent of `print 1 + 2`. To achieve this, we push 1 and 2 to the stack and execute the `ADD` instruction which push the result of summing everything on the stack.

Look in the `code/vm` directory for the full language using this VM.

GOING FURTHER

Building your first language is fun, but it's only the tip of the iceberg. There's so much to discover in that field. Here are a few things I've been playing with in the last years.

HOMOICONICITY

That's the word you want to ostentatiously throw in a geek conversation. While it sounds obscure and complex, it means that the primary representation of your program (the AST) is accessible as a data structure inside the runtime of the language. You can inspect and modify the program as it's running. This gives you godlike powers.

Look in the [Interesting Languages](#) section of the References chapter for the *Io* language and in the [Appendix: Mio, a minimalist homoiconic language](#) of this book for a sample homoiconic language implementing `if` and boolean logic in itself.

SELF-HOSTING

A self-hosting, or metacircular interpreter aims to implement the interpreter in the target language. This is very tedious since you need to implement an interpreter first to run the language, which causes a circular dependency between the two.

[CoffeeScript](#) is a little language that compiles into JavaScript. The CoffeeScript compiler is itself [written in CoffeeScript](#).

[Rubinius](#) is a Ruby implementation that aims to be self-hosted in the future. At the moment, some parts of the runtime are still not written in Ruby.

[PyPy](#) is trying to achieve this in a much simpler way: by using a restrictive subset of the Python language to implement Python itself.

WHAT'S MISSING?

If you're serious about building a real language (real as in production-ready), then you should consider implementing it in a faster and more robust environment. Ruby is nice for quick prototyping but horrible for language implementation.

The two obvious choices are Java on the JVM, which gives you a garbage collector and a nice collection of portable libraries, or C/C++, which gives you total control over what you're doing.

Now get out there and make your own awesome language!

RESOURCES

BOOKS & PAPERS

[Language Implementation Patterns](#), by Terence Parr, from The Programmatic Programmers.

[Smalltalk-80: The Language and its Implementation](#) by Adele Goldberg and al., published by Addison-Wesley, May 1983.

[A No-Frills Introduction to Lua 5.1 VM Instructions](#), by Kein-Hong Man.

[The Implementation of Lua 5.0](#), by Roberto Ierusalimschy et al.

EVENTS

[OOPSLA](#), The International Conference on Object Oriented Programming, Systems, Languages and Applications, is a gathering of many programming language authors.

The [JVM Language Summit](#) is an open technical collaboration among language designers, compiler writers, tool builders, runtime engineers, and VM architects for sharing experiences as creators of programming languages for the JVM.

FORUMS AND BLOGS

[Lambda the Ultimate](#), The Programming Languages Weblog, discuss about new trends, research papers and various programming language topics.

INTERESTING LANGUAGES

[lo](#):

Io is a small, prototype-based programming language. The ideas in Io are mostly inspired by Smalltalk (all values are objects, all messages are dynamic), Self (prototype-based), NewtonScript (differential inheritance), Act1 (actors and futures for concurrency), LISP (code is a runtime inspectable/modifiable tree) and Lua (small, embeddable).

A few things to note about Io. It doesn't have any parser, only a lexer that converts the code to Message objects. This language is Homoiconic.

[Factor](#) is a concatenative programming language where references to dynamically-typed values are passed between words (functions) on a stack.

[Lua](#):

Lua is a powerful, fast, lightweight, embeddable scripting language.

Lua combines simple procedural syntax with powerful data description constructs based on associative arrays and extensible semantics. Lua is dynamically typed, runs by interpreting bytecode for a register-based virtual machine, and has automatic memory management with incremental garbage collection, making it ideal for configuration, scripting, and rapid prototyping.

[tinypy](#) and [tinyrb](#) are both subsets of more complete languages (Python and Ruby respectively) running on Virtual Machines inspired by Lua. Their code is only a few thousand lines long. If you want to introduce a VM in your language design, those are good starting points.

Need inspiration for your awesome language? Check out Wikipedia's programming lists: (List of programming languages)[http://en.wikipedia.org/wiki/List_of_programming_languages], (Esoteric programming languages)[http://en.wikipedia.org/wiki/Esoteric_programming_language]

In addition to the languages we know and use every day (C, C++, Perl, Java, etc.), you'll find many lesser-known languages, many of which are very interesting. You'll even find some esoteric languages such as [Piet](#), a language that is programmed using images that look like abstract art. You'll also find some languages that are voluntarily impossible to use like [Malbolge](#) and [BrainFuck](#) and some amusing languages like [LOLCODE](#), whose sole purpose is to be funny.

While some of these languages aren't practical, they can widen your horizons, and alter your conception of what constitutes a computer language. If you're going to design your own language, that can only be a good thing.

SOLUTIONS TO DO IT YOURSELF

SOLUTIONS TO DO IT YOURSELF I

a. Modify the lexer to parse: `while` condition: ... control structures.

Simply add `while` to the `KEYWORD` array on line 2.

```
1 KEYWORDS = ["def", "class", "if", "else", "true", "false", "nil", "while"]
```

b. Modify the lexer to delimit blocks with `{ ... }` instead of indentation.

Remove all indentation logic and add an `elsif` to parse line breaks.

```
1 class BracketLexer
2   KEYWORDS = ["def", "class", "if", "else", "true", "false", "nil"]
3
4   def tokenize(code)
5     code.chomp!
6     i = 0
7     tokens = []
8
9     while i < code.size
10      chunk = code[i..-1]
11
12      if identifier = chunk[/\A([a-z]\w*)/, 1]
13        if KEYWORDS.include?(identifier)
14          tokens << [identifier.upcase.to_sym, identifier]
15        else
16          tokens << [:IDENTIFIER, identifier]
17        end
18        i += identifier.size
19
20      elsif constant = chunk[/\A([A-Z]\w*)/, 1]
21        tokens << [:CONSTANT, constant]
22        i += constant.size
```

bracket_lexer.rb

```

23
24     elsif number = chunk[/\A([0-9]+)/, 1]
25         tokens << [:NUMBER, number.to_i]
26         i += number.size
27
28     elsif string = chunk[/\A"(.*)""/, 1]
29         tokens << [:STRING, string]
30         i += string.size + 2
31
32     #####
33     # All indentation magic code was removed and only this elsif was added.
34     elsif chunk.match(/\A\n+/)
35         tokens << [:NEWLINE, "\n"]
36         i += 1
37     #####
38
39     elsif chunk.match(/\A /)
40         i += 1
41
42     else
43         value = chunk[0,1]
44         tokens << [value, value]
45         i += 1
46
47     end
48
49 end
50
51 tokens
52 end
53 end

```

SOLUTIONS TO DO IT YOURSELF II

a. Add a rule in the grammar to parse `while` blocks.

This rule is very similar to `If`.

```

1 # At the top add:
2 token WHILE
3
4 # ...
5
6 Expression:
7 # ...
8 | While
9 ;
10
11 # ...
12
13 While:
14 WHILE Expression Block { result = WhileNode.new(val[1], val[2]) }
15 ;

```

And in the `nodes.rb` file, you will need to create the class:

```

1 class WhileNode < Struct.new(:condition, :body); end

```

b. Add a grammar rule to handle the ! unary operators.

Similar to the binary operator. Calling `!x` is like calling `x.!`.

```

1 Operator:
2 # ...
3 | '!' Expression { result = CallNode.new(val[1], val[0], []) }
4 ;

```

SOLUTIONS TO DO IT YOURSELF III

a. Implement inheritance by adding a superclass to each Awesome class.

```

1 class AwesomeClass < AwesomeObject
2   # ...
3
4   def initialize(superclass=nil)

```

```

5     @runtime_methods = {}
6     @runtime_superclass = superclass
7     # ...
8 end
9
10 def lookup(method_name)
11     method = @runtime_methods[method_name]
12     unless method
13         if @runtime_superclass
14             return @runtime_superclass.lookup(method_name)
15         else
16             raise "Method not found: #{method_name}"
17         end
18     end
19     method
20 end
21 end
22
23 # ...
24
25 Runtime["Number"] = AwesomeClass.new(Runtime["Object"])

```

b. Add the method to handle $x + 2$.

```

1 Runtime["Number"].runtime_methods["+"] = proc do |receiver, arguments|
2     result = receiver.ruby_value + arguments.first.ruby_value
3     Runtime["Number"].new_with_value(result)
4 end

```

SOLUTIONS TO DO IT YOURSELF IV

a. Implement the `WhileNode`.

`while` is very similar to `if`.

```

1 class WhileNode
2     def eval(context)
3         while @condition.eval(context).ruby_value

```

```
4     @body.eval(context)
5   end
6 end
7 end
```

APPENDIX: MIO, A MINIMALIST HOMOICONIC LANGUAGE

HOMOICOWHAT?

Homoiconicity is a hard concept to grasp. The best way to understand it fully is to implement it. That is the purpose of this section. It should also give you glimpse at an unconventional language.

We'll build a tiny language called Mio (for mini-lo). It is derived from the [lo language](#). The central component of our language will be messages. Messages are a data type in Mio and also how programs are represented and parsed, thus its homoiconicity. We'll again implement the core of our language in Ruby, but this one will take less than 200 lines of code.

MESSAGES ALL THE WAY DOWN

Like in Awesome, everything is an object in Mio. Additionally, a program being method calls and literals, is simply a series of messages. And messages are separated by spaces not dots, which makes our language looks a lot like plain english.

```
1 object method1 method2(argument)
```

Is the semantic equivalent of the following Ruby code:

```
1 object.method1.method2(argument)
```

THE RUNTIME

Unlike Awesome but like Javascript, Mio is prototype-based. Thus, it doesn't have any classes or instances. We create new objects by cloning existing ones. Objects don't have classes, but prototypes (`protos`), their parent objects.

Mio objects are like dictionaries or hashes (again, much like Javascript). They contain slots in which we can store methods and values such as strings, numbers and other objects.

mio/object.rb

```
1 module Mio
2   class Object
3     attr_accessor :slots, :protos, :value
4
5     def initialize(proto=nil, value=nil)
6       @protos = [proto].compact
7       @value = value
8       @slots = {}
9     end
10
11     # Lookup a slot in the current object and protos.
12     def [](name)
13       return @slots[name] if @slots.key?(name)
14       message = nil
15       @protos.each { |proto| return message if message = proto[name] }
16       raise Mio::Error, "Missing slot: #{name.inspect}"
17     end
18
19     # Set a slot
20     def []=(name, message)
21       @slots[name] = message
22     end
23
24     # The call method is used to eval an object.
25     # By default objects eval to themselves.
26     def call(*)
27       self
28     end
```



```

29
30     def clone(val=nil)
31         val ||= @value && @value.dup rescue TypeError
32         Object.new(self, val)
33     end
34 end
35 end

```

Mio programs are a chain of messages. Each message being a token. The following piece of code:

```

1 "hello" print
2 1 to_s print

```

is parsed as the following chain of messages:

```

1 Message.new('"hello"',
2   Message.new("print",
3     Message.new("\n",
4       Message.new("1",
5         Message.new("to_s",
6           Message.new("print"))))))))

```

Notice line breaks (and dots) are also messages. When executed, they simply reset the receiver of the message.

```

1 self print # <= Line break resets the receiver to self
2 self print # So now it looks as if we're starting a new expression
3           # with the same receiver as before.

```

This results in the same behaviour as in languages such as Awesome, where each line is an expression.

The unification of all types of expression into one data type makes our language extremely easy to parse (see `parse_all` method in the code bellow). Messages

are much like tokens, thus our parsing code will be similar to the one of our lexer in Awesome. We don't even need a grammar with parsing rules!

mio/message.rb

```
1 module Mio
2   # Message is a chain of tokens produced when parsing.
3   #   1 print.
4   # is parsed to:
5   #   Message.new("1",
6   #               Message.new("print"))
7   # You can then +call+ the top level Message to eval it.
8   class Message < Object
9     attr_accessor :next, :name, :args, :line, :cached_value
10
11    def initialize(name, line)
12      @name = name
13      @args = []
14      @line = line
15
16      # Literals are static values, we can eval them right
17      # away and cache the value.
18      @cached_value = case @name
19        when /^#\d+/
20          Lobby["Number"].clone(@name.to_i)
21        when /^"(.*)"$/
22          Lobby["String"].clone($1)
23        end
24
25      @terminator = [".", "\n"].include?(@name)
26
27      super(Lobby["Message"])
28    end
29
30    # Call (eval) the message on the +receiver+.
31    def call(receiver, context=receiver, *args)
32      if @terminator
33        # reset receiver to object at begining of the chain.
34        # eg.:
35        #   hello there. yo
36        #   ^           ^__ "." resets back to the receiver here
37        #   \_____ /
38        value = context
```

```

39     elsif @cached_value
40         # We already got the value
41         value = @cached_value
42     else
43         # Lookup the slot on the receiver
44         slot = receiver[name]
45
46         # Eval the object in the slot
47         value = slot.call(receiver, context, *@args)
48     end
49
50     # Pass to next message if some
51     if @next
52         @next.call(value, context)
53     else
54         value
55     end
56 rescue Mio::Error => e
57     # Keep track of the message that caused the error to output
58     # line number and such.
59     e.current_message ||= self
60     raise
61 end
62
63 def to_s(level=0)
64     s = " " * level
65     s << "<Message @name=#{@name}"
66     s << ", @args=" + @args.inspect unless @args.empty?
67     s << ", @next=\n" + @next.to_s(level + 1) if @next
68     s + ">"
69 end
70
71 # Parse a string into a chain of messages
72 def self.parse(code)
73     parse_all(code, 1).last
74 end
75
76 private
77 def self.parse_all(code, line)
78     code = code.strip
79     i = 0
80     message = nil

```

```

81     messages = []
82
83     # Marrrvelous parsing code!
84     while i < code.size
85         case code[i..-1]
86             when /\A("[^"]*"\/, # string
87                 /\A(\d+\/, # number
88                 /\A(\.)+\/, # dot
89                 /\A(\n)+\/, # line break
90                 /\A(\w+\/ # name
91             m = Message.new($1, line)
92             if messages.empty?
93                 messages << m
94             else
95                 message.next = m
96             end
97             line += $1.count("\n")
98             message = m
99             i += $1.size - 1
100         when /\A(\(\s*)\/ # arguments
101             start = i + $1.size
102             level = 1
103             while level > 0 && i < code.size
104                 i += 1
105                 level += 1 if code[i] == ?\ (
106                 level -= 1 if code[i] == ?\ )
107             end
108             line += $1.count("\n")
109             code_chunk = code[start..i-1]
110             message.args = parse_all(code_chunk, line)
111             line += code_chunk.count("\n")
112         when /\A,(\s*)\/
113             line += $1.count("\n")
114             messages.concat parse_all(code[i+1..-1], line)
115             break
116         when /\A(\s+\/, # ignore whitespace
117             /\A(#.*$\/ # ignore comments
118             line += $1.count("\n")
119             i += $1.size - 1
120         else
121             raise "Unknown char #{code[i].inspect} at line #{line}"
122         end

```

```

123         i += 1
124     end
125     messages
126 end
127 end
128 end

```

The only missing part of our language at this point is a method. This will allow us to store a block of code and execute it later in its original context and on the receiver.

But, there will be one special thing about our method arguments. They won't be implicitly evaluated. For example, calling `method(x)` won't evaluate `x` when calling the method, it will pass it as a message. This is called lazy evaluation. It will allow us to implement control structure right from inside our language. When an argument needs to be evaluated, we do so explicitly by calling the method `eval_arg(arg_index)`.

```

1  module Mio
2    class Method < Object
3      def initialize(context, message)
4        @definition_context = context
5        @message = message
6        super(Lobby["Method"])
7      end
8
9      def call(receiver, calling_context, *args)
10         # Woo... lots of contexts here... lets clear that up:
11         #   @definition_context: where the method was defined
12         #   calling_context: where the method was called
13         #   method_context: where the method body (message) is executing
14         method_context = @definition_context.clone
15         method_context["self"] = receiver
16         method_context["arguments"] = Lobby["List"].clone(args)
17         # Note: no argument is evaluated here. Our little language only has lazy argument
18         # evaluation. If you pass args to a method, you have to eval them explicitly,
19         # using the following method.

```

mio/method.rb

```

20     method_context["eval_arg"] = proc do |receiver, context, at|
21         (args[at.call(context).value] || Lobby["nil"]).call(calling_context)
22     end
23     @message.call(method_context)
24 end
25 end
26 end

```

Now that we have all the objects in place we're ready to bootstrap our runtime.

Our Awesome language had a `Context` object, which served as the environment of execution. In Mio, we'll simply use an object as the context of evaluation. Local variables will be stored in the slots of that object. The root object is called the Lobby. Because ... it's where all the objects meet, in the lobby. (Actually, the term is taken from Io.)

```

1  module Mio
2      # Bootstrap
3      object = Object.new
4
5      object["clone"] = proc { |receiver, context| receiver.clone }
6      object["set_slot"] = proc do |receiver, context, name, value|
7          receiver[name.call(context).value] = value.call(context)
8      end
9      object["print"] = proc do |receiver, context|
10         puts receiver.value
11         Lobby["nil"]
12     end
13
14     # Introducing the Lobby! Where all the fantastic objects live and also the root context
15     # of evaluation.
16     Lobby = object.clone
17
18     Lobby["Lobby"] = Lobby
19     Lobby["Object"] = object
20     Lobby["nil"] = object.clone(nil)
21     Lobby["true"] = object.clone(true)
22     Lobby["false"] = object.clone(false)

```

mio/bootstrap.rb

```

23 Lobby["Number"] = object.clone(0)
24 Lobby["String"] = object.clone("")
25 Lobby["List"]   = object.clone([])
26 Lobby["Message"] = object.clone
27 Lobby["Method"] = object.clone
28
29 # The method we'll use to define methods.
30 Lobby["method"] = proc { |receiver, context, message| Method.new(context, message) }
31 end

```

IMPLEMENTING MIO IN MIO

This is all we need to start implementing our language in itself.

First, here's what we're already able to do: cloning objects, setting and getting slot values.

```

1 # Create a new object, by cloning the master Object
2 set_slot("dude", Object clone)
3 # Set a slot on it
4 dude set_slot("name", "Bob")
5 # Call the slot to retrieve it's value
6 dude name print
7 # => Bob
8
9 # Define a method
10 dude set_slot("say_name", method(
11   # Print unevaluated arguments (messages)
12   arguments print
13   # => <Message @name="hello...">
14
15   # Eval the first argument
16   eval_arg(0) print
17   # => hello...
18
19   # Access the receiver via `self`
20   self name print
21   # => Bob
22 ))

```

test/mio/oop.mio

23

```
24 # Call that method
25 dude say_name("hello...")
```

Here's where the lazy argument evaluation comes in. We're able to implement the `and` and `or` operators from inside our language.

```
1 # An object is always truish
2
3 Object set_slot("and", method(
4   eval_arg(0)
5 ))
6 Object set_slot("or", method(
7   self
8 ))
9
10 # ... except nil and false which are false
11
12 nil set_slot("and", nil)
13 nil set_slot("or", method(
14   eval_arg(0)
15 ))
16
17 false set_slot("and", false)
18 false set_slot("or", method(
19   eval_arg(0)
20 ))
```

mio/boolean.mio

```
1 "yo" or("hi") print
2 # => yo
3
4 nil or("hi") print
5 # => hi
6
7 "yo" and("hi") print
8 # => hi
9
10 1 and(2 or(3)) print
11 # => 2
```

test/mio/boolean.mio

Using those two operators, we can implement `if`.

mio/if.mio

```
1 # Implement if using boolean logic
2
3 set_slot("if", method(
4   # eval condition
5   set_slot("condition", eval_arg(0))
6   condition and( # if true
7     eval_arg(1)
8   )
9   condition or( # if false (else)
10    eval_arg(2)
11  )
12 ))
```

And now... holy magical pony!

test/mio/if.mio

```
1 if(true,
2   "condition is true" print,
3   # else
4   "nope" print
5 )
6 # => condition is true
7
8 if(false,
9   "nope" print,
10  # else
11  "condition is false" print
12 )
13 # => condition is false
```

`if` defined from inside our language!

BUT IT'S UGLY

All right... it's working, but the syntax is not as nice as Awesome. An addition would be written as: `1 + (2)` for example, and we need to use `set_slot` for assignment, nothing to impress your friends and foes.

To solve this problem, we can again borrow from Io and implement operator shuffling. This simply means reordering operators. During the parsing phase, we would turn `1 + 2` into `1 + (2)`. Same goes for ternary operators such as assignment. `x = 1` would be rewritten as `= (x, 1)`. This introduces syntactic sugar into our language without impacting its homoiconicity and awesomeness.

You can find all the source code for Mio under the `code/mio` directory and run its unit tests with the command: `ruby -Itest test/mio_test.rb`.

FAREWELL!

That is all for now. I hope you enjoyed my book!

If you find an error or have a comment or suggestion, please send me an email at macournoyer@gmail.com.

If you end up creating a programming language let me know, I'd love to see it!

Thanks for reading.

- *Marc*